

Buyer's Guide: Computer-Based Training for Enterprise IT

# Windows IT Pro

A PENTON PUBLICATION

JUNE 2012 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

## Getting Started with **PowerShell**

Windows Server 2012  
Storage Spaces

**Exchange Server 2010**  
Corrupted Items and Mailbox Moves

Automate with Windows Task Scheduler

**BitLocker in Windows 8**  
Enhance Data Protection

Network Monitoring with  
System Center 2012 Operations Manager

Pros and Cons of SharePoint in the Cloud

# We've Mapped the DNA of Business Continuity.

## Double-Take® 6.0

Our category-leading, platform-agnostic solutions for replication, migration, availability and disaster recovery offer six distinct advantages:

**Simplicity**

**Plug & Play Protection**

**Real-Time Replication**

**No-Downtime Migration**

**Platform Independence**

**Unified Console**

[visionsolutions.com](http://visionsolutions.com)



© Copyright 2012, Vision Solutions, Inc. All rights reserved. Vision Solutions® and Double-Take® are registered trademarks of Vision Solutions, Inc.



# BYOD – Risky Business?



The Bring Your Own Device trend is escalating. Unprecedented numbers of employees are bringing their Macs, iPads and iPhones, connecting them to corporate email and network services. How can you meet the associated security and compliance challenges? With Centrify, these Mac and Mobile devices can be associated with the user account in Active Directory. No additional management tools or infrastructure needed. Even use Windows Group Policy to enforce security settings. Best of all: it's FREE.

**Now safely say yes to Macs and iPads. With Centrify, you can centrally secure and manage mobile devices using Active Directory.**

**Download FREE Centrify Express for Mobile.**

**Download Now** ►



**Visit us at TechEd 2012, Booth 239 and get a free copy of Express!**



# COVER STORY ▼

47

## Getting Started with Windows PowerShell — Don Jones

Why are so many administrators afraid of little old PowerShell? Probably because they don't understand it. Don't let common misconceptions stand in your way—get to know the power of PowerShell, and increase your value in the workforce.

### Features

#### 57 Windows Server 2012 Storage Spaces

John Savill

#### 69 Corrupted Items and Mailbox Moves in Exchange Server 2010

Tony Redmond

#### 77 Task Scheduler: All Grown Up

Rob Gravelle

#### 93 BitLocker in Windows 8

Jan De Clercq

#### 103 Integrated Network Monitoring in System Center 2012 Operations Manager

John Joyner

#### 116 SharePoint in the Cloud

Michael Noel

### Interact

#### 39 Ask the Experts

### In Every Issue

#### 10 IT Community Forum

#### 165 Ctrl+Alt+Del

#### 167 Advertiser Directory

#### 167 Directory of Services

#### 167 Vendor Directory

### Chat with Us



Facebook



Twitter



LinkedIn



# Columns

7

IT Pro Perspectives

## The Emergence of the Hybrid Cloud

Michael Otey



13

Need to Know

## Office 15

Paul Thurrott



19

Windows Power Tools

## Searching and Managing Active Directory Groups with PowerShell

Mark Minasi



23

Top 10

## Internet Explorer 9 Tips

Michael Otey



27

Enterprise Identity

## How Windows Server 2012 Improves Active Directory Disaster Recovery

Sean Deuby



32

What Would Microsoft Support Do?

## The Road Warrior's Laptop Build Guide

Richmond V. Baker, Roger Osborne, Joe Quint, and Hollis Williams



# Products

## 127 New & Improved

## 131 Paul's Picks

Paul Thurrott

## 132 XenDesktop 5.6

Michael Otey

## 139 PowerBroker Desktops, Windows Edition

Eric B. Rux

## 143 SharePoint Management and Diagnostic Tool Roundup

Russell Smith

## 153 Computer-Based Training for Enterprise IT Systems

Tony Bieda

## 160 Industry Bytes

## Editorial

Editorial Director:

Megan Keller

Editor in Chief:

Amy Eisenberg

Senior Technical Director:

Michael Otey

Technical Director:

Sean Deuby

Senior Technical Analyst:

Paul Thurrott

Industry News Analyst:

Jeff James

Custom Group Editorial Director:

Dave Bernard

Exchange & Outlook:

Brian Winstead

Systems Management,

Networking, Hardware:

Jason Bovberg

Security, Virtualization:

Jeff James

SharePoint, Active Directory:

Caroline Marwitz

SQL Server, Developer Content:

Megan Keller

Managing Editor:

Lavon Peters

Editorial SEO Specialist:

Jayleen Heft

Assistant Editor:

Blair Greenwood

## Senior Contributing Editors

David Chernicoff, Mark Minasi,

Paul Robichaux, Mark Russinovich,

John Savill

## Contributing Editors

Alex K. Angelopoulos, Michael Dragone,

Jeff Felling, Brett Hill, Dan Holme,

Darren Mar-Elia, Tony Redmond,

Eric B. Rux, William Sheldon,

Curt Spanburgh, Bill Stewart, Orin Thomas,

Douglas Toombs, Ethan Wilansky

## Art & Production

Production Director: Linda Kirchgesler

Senior Graphic Designer: Matt Wiebe

## Advertising Sales

Publisher: Peg Miller

Key Account Director:

Chrissy Ferraro • 970-203-2883

Account Executives:

Barbara Ritter • 858-367-8058

Cass Schulz • 858-357-7649

## Client Project Managers

Michelle Andrews • 970-613-4964

Kim Eck • 970-203-2953

Ad Production Supervisor:

Glenda Vaught

## Marketing & Circulation

Customer Service

Senior Director, Marketing Analytics:

Tricia Syed

Online Sales Development Director:

Amanda Phillips • 970-203-2806

## Technology Group

Senior VP, Penton Media Technology Group:

Kim Paulsen

## Corporate

Chief Executive Officer:

David Kieselstein

Chief Financial Officer/Executive Vice

President: Nicola Allais



## List Rentals

MeritDirect

333 Westchester Avenue, White Plains, NY

## Reprints

Reprint Sales:

Wright's Media • 877-652-5295

*Windows IT Pro*, June 2012, Issue no. 214, ISSN 1552-3136.

*Windows IT Pro* is published monthly by Penton Media, Inc. Copyright ©2012 Penton Media, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any way without the written consent of Penton Media, Inc.

*Windows IT Pro*, 748 Whalers Way, Fort Collins, CO 80525, 800-621-1544 or 970-663-4700. Customer Service: 800-793-5697.

We welcome your comments and suggestions about the content of *Windows IT Pro*. We reserve the right to edit all submissions. Letters should include your name and address. Please direct all letters to [letters@windowsitpro.com](mailto:letters@windowsitpro.com). IT pros interested in writing for *Windows IT Pro* can submit articles to [articles@windowsitpro.com](mailto:articles@windowsitpro.com).

Program Code: Unless otherwise noted, all programming code in this issue is ©2012, Penton Media, Inc., all rights reserved. These programs may not be reproduced or distributed in any form without permission in writing from the publisher. It is the reader's responsibility to ensure procedures and techniques used from this publication are accurate and appropriate for the user's installation. No warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server® are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries and are used by Penton Media, Inc., under license from owner. *Windows IT Pro* is an independent publication not affiliated with Microsoft Corporation. Microsoft Corporation is not responsible in any way for the editorial policy or other contents of the publication.

# Windows IT Pro

# SPECOPS:software

Migrate to Windows 7.  
Faster. Much, much faster.



Download your free trial version today.



Meet us at TechEd North America, booth 606  
for a demonstration of Specops Deploy and  
your chance to Win a Trip for Two to the  
to the Daytona 500.

[specopssoft.com](http://specopssoft.com)

Innovation & Simplicity. That's Specops.





# The Emergence of the Hybrid Cloud

Unless you've been hiding under a rock for the past couple of years, you certainly know that the cloud has been a major technological push by almost all of the major IT vendors. No doubt this move is in part fueled by the ubiquitous nature of the Internet. But it has also been growing partly in response to the challenging economic conditions of the past several years. At its essence, cloud technology promises to save organizations money by allowing them to buy only the level of IT services that they need without the added expense of the IT infrastructure required to run it.

While there's little doubt that cloud will be an important IT technology going forward, it's also clear that businesses haven't been stumbling all over themselves to buy into the technology. A number of hurdles and considerations such as security, performance, and the need to change applications have stopped many businesses from moving to the early cloud technologies. That said, it's also certain that cloud technology has been evolving rapidly.

At the recent [Microsoft Management Summit \(MMS\) 2012](#) in Las Vegas, it became clear that the newest evolution of cloud technology is the hybrid cloud. The hybrid cloud combines the public cloud with your own private cloud infrastructure. For example, you might implement your database services on a private cloud that's running on your own internal infrastructure in conjunction with an application layer that's running on a public cloud infrastructure.

Using this type of hybrid cloud infrastructure over a pure public cloud or private cloud infrastructure offers some significant advantages. The application workload has the potential for the most volatility,



## Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).



Email

and running the application layer on the public cloud infrastructure enables the application front end to be easily scaled up and down to meet changing demands. The back-end database workload tends to be more predictable, and running the database tier on your own private cloud infrastructure lets you have complete control over your company's data.

**Video**

Michael Otey on the  
significance of the  
hybrid cloud

---



Under this scenario, security and backups are completely within the control of your organization. Further, in many countries, regulatory requirements mandate that a business's data remain within the boundaries of that country. Keeping the database on premises allows the organization to have complete control over where the data is located. That's not case if the data were stored in the public cloud where each individual company or client has no control over the cloud vendor's infrastructure and where their storage is located.

One example of a technology that can help implement this type of hybrid cloud model can be found in the [HP Database Consolidation Appliance](#). The HP appliance is a private cloud infrastructure that's

specially designed to support multiple virtualized database workloads. The appliance functions as a private cloud for your data tier, enabling both high performance through its optimized internal storage and networking configurations as well as high availability through its built-in failover clustering services. This type of database private cloud can then be used to support applications running either in the public cloud or on your on-premises infrastructure.

The hybrid cloud has its considerations as well. Probably the biggest concern is the last-mile problem. The network connection between the private cloud infrastructure and public cloud represents a potential single point of failure and, as such, requires redundancy to ensure that your public cloud applications can connect to your private cloud database resources.

Managing multiple clouds and multiple cloud technologies adds complexity and requires management tools that can span both technologies. Today, this typically means that you'll need to use private and public cloud technologies that come from the same vendor. For example, with the previous HP database appliance example, you would need to use Microsoft System Center to manage the database private cloud as well as the public cloud application running in Windows Azure.

Cloud technologies are quickly evolving, and new cloud implementations such as the hybrid cloud are emerging to address the problems and limitations that were found in the early cloud implementations. The hybrid cloud solves a lot of the problems that hindered the adoption of early cloud technologies by combining the best attributes of the public cloud and the private cloud. The public cloud brings built-in scalability and availability to the table; the private cloud allows you to have complete control over your corporate data. ■

InstantDoc ID 142944



# Letters

[letters@windowsitpro.com](mailto:letters@windowsitpro.com)

## Windows 8 Equals More Clicks

As a software developer, I once wrote down a mantra: “Fewer clicks equals fewer tech calls.” That statement is a fact. I can prove it by looking up any Dell or HP or Toshiba call log. Which issues get the most calls? They generally involve the tasks with the most clicks.

My car's steering wheel has buttons to control the radio now. But the actual radio still has knobs, too. The world likes tradition; the world likes comfort. Don't try to push your idea on the masses. Let them migrate at will. You always need to leave the old layout intact, with options to turn on the new layout for those who are interested.

Most people aren't geeks. They don't want to learn new systems. They just want to get things done. Why would we change how we start our cars or unlock our doors, if we're not reducing the time it takes to complete the task? That's the way to lose customers.

OSs should have skins. That way, you could use the same skin for the rest of your life without ever having to relearn basic tasks. The general population doesn't like change, especially when it's something personal that they use many times a day. Change is for a small group of individuals. Geeks just need options in Control Panel for geekifying their computers, but those options should always be turned off by default. The geek will know where to find them. The average person just isn't interested enough.

Aero is ridiculous—nobody needs it or wants it. It serves no purpose except to confuse people. Oh, and the Ribbon idea is terrible,

## Send Your Comments

*Windows IT Pro* welcomes feedback about the magazine. Send comments, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

Comments



inefficient, impossible to train on, and completely non-intuitive. Menus, menus, menus...they've worked in restaurants for millennia. Menus are logical. I always know that File is first, Edit is second, View is third, and Help is last. So simply standardized, and perfectly intuitive. And Windows Vista? Vista is actually a good OS that kept the user experience intact. I will stick to that for as long as possible.

—John Becker

## Carrier Bloatware

I want to thank B. K. Winstead for his article "[Carrier Bloatware: The Android Plague](#)." I learned a hard lesson after I bought my Droid Charge: Always check how much run-time memory a phone has before buying it. My phone, because of the Samsung and Verizon bloatware, is only good for people who don't buy apps or those that root the phone and remove the bloatware. I chose the latter.

—Paul Springer

*Thanks for writing! It's always great to hear from readers. And thanks for the app suggestions you left in the article comments—"Startup Manager (to limit what runs), Advanced Task Manager (to kill tasks that shouldn't be running), and SystemPanel (the best utility app ever!)." I've heard of them but haven't tried them. I probably should. Honestly, as I wrote that article, I had to rein in my desire to rail against the carriers. I firmly believe something needs to be done about this problem, but the carriers have all the power as things stand. Perhaps a groundswell of consumer anger will occur eventually.*

—B. K. Winstead

## Another Free Security Tool

I reviewed Jeff James's list of "[13 Free Security Tools and Resources](#)" and wanted to add another great tool to it. [Paessler Router Traffic Grapher](#) (PRTG) is a free tool available on CNET. Aside from the usual SNMP stats and trigger alerts that most network tools offer, the real

gem of this tool is its ability to view network traffic by protocol via graphs. You can add your own *channel* and include it in the graph. A channel can be an IP address or protocol. I use this tool to monitor my Internet link, monitor the LAN port of my firewall, and monitor any internal computer device I want—all at the same time. The product is solid and hasn't crashed once in the 5 years that I've used it.

—Pierre Massé

## Server Core in Server 2012

I remember reading Mark Minasi's article "[Sampling Server Core](#)," in which he writes, "[I] fear that Server Core might not be widely used, because, well, while a stripped-down server might be attractive in many ways, most of us expect a stripped-down price as well. After all, I love my Honda Insight hybrid with its three-cylinder engine and 67mpg fuel efficiency, but I wouldn't have purchased it if it had cost as much as a Jaguar." At the time, Mark's opinion made perfect sense to me. I found his comment clever, fair, and insightful. But with Windows Server 2012 just around the corner, I see things differently. Server 2012 has Server Core as an option that you can turn on and off at will. So, in hindsight, I realize that Microsoft was right all along to offer Windows Server and Server Core for the same price, since Microsoft really meant Server Core to be a mode of operation that could be enabled and disabled at will and not a different product.

—Dimitrios Kalemis

*So true, Dimitrios, but—and this is the important point—we don't know the price yet. You and I are playing with the Datacenter version of Server 2012. For all I (or you) know, you'll have to have the expensive Enterprise version to get these wonderful features. Hey, I hope it isn't true, but it wouldn't surprise me all that much. It would, however, sadden me. Thanks for reading!* ■

—Mark Minasi

InstantDoc ID 142961



# Office 15

**T**he next several months should prove to be monumental from a Microsoft platforms perspective. But it's not just Windows 8, Windows Server 2012 (formerly code-named Windows Server 8), and Windows Phone 8 that promise to recast Microsoft's core platforms for the future. The software giant is also working on its most massive upgrade yet to the Office family of products. And if all goes well, Microsoft will for the first time upgrade virtually all of the products and services within the Office sphere at the same time. It's an audacious plan that could unravel for any number of reasons.

This next Office wave, dubbed Office 15, encompasses not just the traditional, PC-based productivity suite, but also the Office servers (Exchange, SharePoint, Lync), hosted online services (Office 365), web-based Office Web Apps (for SkyDrive and SharePoint), various mobile apps, and more. As with the Windows "better together" strategy, Microsoft feels it can make a great case for why these solutions, although great on their own, work better when upgraded in tandem. And although staggered releases have been generally trouble-free in the past (we'll forget about that version of Outlook that wasn't full-featured until a delayed Exchange Server shipped), the plan for this round isn't so much wave as it is tsunami.

With Office 15, the venerable productivity platform evolves yet again for a new age. Office, now the major part of Microsoft's biggest business unit (by revenues), is at a crossroads where the future direction will dramatically affect not just Office but also Microsoft. With Windows revenues decreasing, and the Windows 8 future cloudy thanks to the rise of post-PC device platforms such as the Apple iPad, the question is simple: Should Microsoft push Office applications and functionality on alternative platforms at the expense of Windows? Or should the company keep the best Office capabilities in-house and only on Windows?



## Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for Windows IT Pro UPDATE, and a daily Windows news and information newsletter called WinInfo Daily UPDATE.



**Email**



**Twitter**



**Website**

We'll know which way the pendulum swings soon enough, if not at the Microsoft TechEd 2012 show, then by the fall launch window for Windows 8, Windows Server 2012, and Windows Phone 8. But early indications are that Microsoft is making a bold push with Office toward alternative platforms, starting, if my sources are correct, with Microsoft Outlook, which could be the next major Office mobile app for iPad and Android. Beyond that are plans for mobile versions of Office for iOS (iPhone/iPad) and Android that will rival the Office Mobile versions on Windows Phone 8 and Windows 8 on ARM (Windows on ARM—WOA).

I've been arguing for exactly this strategy for years, and although Windows boss Steven Sinofsky probably does have the clout to push Microsoft back from the brink to focus on a Windows-first strategy, it is in fact his former division, Office, that should be pulling the strings. Put another way, Microsoft's future growth could rely less on its OSs and more on heterogeneous platforms such as Office, System Center and Windows Intune device and PC management, entertainment properties such as Xbox, and other solutions that can integrate well with other systems.

Regardless of this potential future, we do know that Microsoft will deliver what we currently call Office 15. And we know that this traditional software will be supported by a major upgrade to the Office Web Apps and to servers such as Exchange and SharePoint. We'll cover each of these Office 15 initiatives in coming months, but for now let's focus on the Office 15 productivity suite itself.

## Office 15: The Mile-High View

In keeping with the move to electronically delivered software, Office 15 will be made available as a service as well as in traditional software packaging. The service-based Office version uses the Microsoft Application Virtualization (App-V) technologies to deliver the full Office productivity suite in a form that can be more easily managed centrally. Using Microsoft's Click-to-Run technology, you will also be able to

stream Office 15 to users' desktops during installation, so users can be accessing Office before the full suite is even downloaded.

Application virtualization offers several benefits over traditionally deployed software, of course. And although this might seem minor, it means that Office 15 can be installed side by side with a previous Office version, which is a key advantage for those evaluating the new suite. Previously, it was difficult and tedious to run different Office versions side by side or, in the case of applications such as Outlook, impossible to do so.

From an applications perspective, Office 15 isn't blazing the way with new applications, but the suite is so full-featured it doesn't really have any obvious functional holes at this point. Instead, we see an evolution of existing capabilities from application to application and the formalization of a new design language, if you will, that's equally at home on both traditional PCs and the touch-based Windows devices that are expected to become increasingly popular.

Key to this dual-use nature is a new, washed-out user experience that pervades across the app, with merely color-coded accents to break the monotony. According to Microsoft, this is an attempt to highlight the content you're working on and de-emphasize the surrounding UI "chrome." It's a Metro-based design principle, similar to the work Microsoft has previously done with Internet Explorer 9 and 10, Windows Phone, and, of course, Windows 8.

These new Metro-style user experiences also provide a minimized ribbon for a cleaner look and a new Touch Mode that enlarges and spaces out UI elements—buttons, menus, and other controls—making them easier to tap with a finger. This effect is a bit too subtle, and, for the majority of users on traditional PCs, somewhat superfluous.

Each of the applications also provides a new Start Experience, which should be a boon to less-experienced users. (And if you're familiar with the Mac versions of Office, it will be immediately recognizable as an interface that debuted there first.) This experience basically provides a visual grid of prebuilt templates so you can create a



nicely designed document (or whatever) from the get-go. Power users will be excited to discover they can disable this screen.

Finally, with this version, Office is embracing the cloud and offering access to SharePoint- and SkyDrive-based document repositories in a far more seamless fashion than did previous versions. For many, the notion of locally stored documents, locked to a single hard drive on a single PC, will disappear.

## **Excel 15**

Microsoft's earliest Office application picks up a new Quick Analysis Lens to more quickly and easily create visual representations of your data. The application also recommends Charts and PivotTables that are most appropriate for the selected data, rather than let you muddle around through the many available templates.

## **Word 15**

The world's most popular word processor picks up a new Read Mode that uses those Metro ideals to present documents—including, for the first time, PDF documents—in a nice display that automatically reflows, using columns depending on the screen width and orientation. Like many Office 15 applications, Word also features a handy bookmark feature called Resume Reading that helps you pick up document re-edits exactly where you were the last time you used the application. The Navigation pane has been improved to be actually usable in this release, and a new Present Online feature, similar to that in PowerPoint, lets you share documents with others through a browser, even if they don't have Word installed.

## **PowerPoint 15**

Microsoft's presentation package finally defaults to a widescreen 16:9 format, though it will give businesses stuck on the old 4:3 VGA format fits until they figure out how to switch back. New Visual Basic-like positioning guides make it easier than ever to line up items on a slide,

and a dramatically improved Presenter View provides a better-than-ever two-screen experience. PowerPoint provides the Resume Reading feature, too, so you can pick up editing where you left off.

## OneNote 15

OneNote came of age in Office 2010, thanks to two factors: First, OneNote was for the first time included by default with all mainstream versions of the suite. And second, this was when OneNote shed its reliance on PC-based notebooks and opened up to the cloud, a scenario that's now replicating across the other applications in Office 15. To follow up that success, Office 15 provides improved digital ink support to take advantage of the coming generation of new Tablet PCs and other touch devices, better tables support, better cloud-to-PC syncing, and, perhaps most crucially, an ever-expanding set of OneNote apps for mobile platforms including Windows Phone, iPhone, iPad, Android phones, and Nokia Symbian Belle handsets.

## Outlook 15

Hundreds of millions of users live in Outlook each day, making it the center of their professional and personal lives. And with Outlook 15, we're getting a refined user experience with better navigation between the email, calendar, people, and tasks modules, and a new Peeks feature for quickly viewing information about your schedule, a person, a task, and other objects without leaving the current view and navigating to the relevant module. It features in-line replies, a new weather bar (in Calendar only, curiously), and finally integrates correctly with multiple email sources, including Hotmail (without requiring an add-on). Speaking of add-ons, remember the Social Connector from Outlook 2010? That's integrated as well.

## Extensibility

Office 15 is also providing an interesting new extensibility platform, code-named Agave, which will work with both traditional, PC-based

versions of Office and the Office Web Apps and Office servers. Agave provides what Microsoft calls “web-powered experiences,” using a web extension model that uses web standards behind the scenes. Developers will be able to publicize their new add-ons using an Office Marketplace that will be accessible from within the applications and can be used by corporations to deliver secure, private solutions to managed users.

## Timing and Expectations

Microsoft is planning to complete development of all Office 15 products and services in November 2012, though a leaked road map suggests that they won’t actually ship to customers until very early 2013. Regardless of the actual ship date, I think it’s highly probable that one or more of the solutions will slip, and as a result, Microsoft will ship these products over a tight period of time, nearly simultaneously rather than actually simultaneously. On a related note, you might recall that Microsoft is bundling a subset of the full Office suite—the applications Word, Excel, PowerPoint, and OneNote—with WOA, the Windows 8 version aimed at iPad-like, ARM-based tablets and other touch devices. The timing of this release is tied to that of Windows 8, and it will likely ship before the other Office 15 wave products and services.

## What’s Left

The big question, of course, is whether Office 15 offers a compelling upgrade. For those on Office 2010, I’d have to say no. But a huge percentage of Office customers are on older, less capable versions of the suite, and Office 15 is a sizable improvement over both Office 2003 and 2007. If you’re still using those products, you’ll want to evaluate Office 15’s public beta as soon as possible.

Office 15 is compelling and forward-leaning. Microsoft’s embrace of cloud computing and virtualization technologies, in particular, is well done and makes supporting and using this suite easier than ever. That’s true regardless of the Office version you’re currently using. ■

InstantDoc ID: 142847

# Searching and Managing Active Directory Groups with PowerShell

Retrieve sets of users based on group membership

In this series of columns about Windows Server 2008 R2's Active Directory (AD)-related PowerShell cmdlets, I've shown you how to use the *get-aduser* and *search-adaccount* PowerShell cmdlets to extract subsets of your AD user accounts based on a range of criteria, but I haven't yet considered one commonly needed criterion: group membership. This month, I'll use *get-adgroupmember* and some of its related cmdlets to show you how to retrieve sets of users based on group membership.

*Get-adgroupmember's* syntax is pretty simple in its basic form:

```
get-adgroupmember groupname
```

Thus, to view all the members in a group named *folks*, you'd type

```
get-adgroupmember folks
```

To unlock the accounts of all members of the *folks* group, you could feed the output of the *get-adgroupmember* command to another AD cmdlet, *enable-adaccount*:

```
get-adgroupmember folks|enable-adaccount
```



## Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

As you’ve probably guessed, *enable-adaccount* simply enables disabled accounts. (Who says PowerShell is hard to understand?) Also, I haven’t yet discussed the “|” (or pipeline) character in detail, but in short, its job is to take the output of one command and make it the input of another command. That’s pretty neat, because—trust me—doing what that little “one liner” accomplishes using VBScript and the Active Directory Scripting Interface (ADSI) would translate to about 50-75 lines of code, a lot of testing, and a couple days’ work for someone who doesn’t do it all the time. In contrast, I created my PowerShell example in about 3 minutes.

If you’d like to create a few groups to test *get-adgroupmember*, PowerShell has a few cmdlets to create and populate groups. For example, suppose you have three user accounts named *user1*, *user2*, and *user3*. You can quickly generate them with this command:

```
for /l %i in (1 1 3) do (net user user%i Passw0rd /add)
```

(By the way, don’t type that into the PowerShell command prompt; use an old-style Windows command prompt.) Next, create a pair of nested global groups into which you’ll place the users. You’ll create a global group named *manyfolks* and another called *folks*, then put the *folks* group into the *manyfolks* group. PowerShell’s group-creating cmdlet is called *new-adgroup*, which looks like

```
new-adgroup name scope-of-group
```

The group’s “scope” is either a numeric or textual value, and the cmdlet accepts 0 or *domainlocal* for a domain local group, 1 or *global* for a global group, and 2 or *universal* for a universal group. (In case you were thinking that you would have called it a group *type*, *new-adgroup* also has a group type, but it refers to whether a group is a security or distribution group.) The following two commands will create *folks* and *manyfolks*:



```
new-adgroup manyfolks global
new-adgroup folks global
```

Next, you'll want to nest the *folks* group in the *manyfolks* group. You can add either a group or a user to a group with the same command, *add-adgroupmember*. In its most common usage, it looks like

```
add-adgroupmember groupname newmember1,newmember2,newmember3 ...
```

To put the *folks* group and the *user1* account directly in *manyfolks*, type

```
add-adgroupmember manyfolks user1,folks
```

And in case you can't tell, the list *user1,folks* contains no spaces. You put *user2* and *user3* into *folks* like so:

```
add-adgroupmember folks user2,user3
```

With the stage set, you can put *get-adgroupmember* through its paces. If you type

```
get-adgroupmember manyfolks | ft samaccountname
```

you'll see output like

```
samaccountname
-----
user1
folks
```

That's technically correct, but probably not what you wanted. Typically, you'd like to see the users in a nested group regardless of how deeply nested their group is, and you can do that by

adding the *-recursive* switch (which conveniently shortens to just *-r*) like so:

```
get-adgroupmember manyfolks -r | ft samaccountname
```

You'll then get output like

```
samaccountname
-----
user1
user2
user3
```

When instructed to use *-recursive*, *get-adgroupmember* examines every nested group, pulling out just the non-container (i.e., user or computer) objects.

Did you notice something strange about the *add-adgroupmember* and *new-adgroup* commands? If not, look at their syntax again. You'll see that they seem to use two non-named parameters each. For example, a more verbose version of the first *new-adgroup* command would be

```
new-adgroup -name manyfolks -groupscope global
```

Why didn't you need to name the *-name* and *-groupscope* parameters? Because PowerShell lets its cmdlet authors give us the option to forgo the *-whatever* parameters, allowing us to just type some set of parameters in the cmdlet invocation in a particular order, making them "positional" rather than "named" parameters. How did I know which parameters could be positional? I read the Help text, which was just a bit cryptic—but I'll decrypt it for you next month. ■

InstantDoc ID 142824

# Internet Explorer 9 Tips

Increase your daily productivity by learning to pin a website to your taskbar, quickly open a new tab, and other shortcuts in Microsoft's IE 9 browser

**A**lthough many users have moved over to Chrome or Firefox, Internet Explorer (IE) is still the most widely used browser—by far. According to NetMarketShare, Internet Explorer holds about 53.83 percent of the browser market, followed by Firefox with 22.87 percent and Chrome at 13.19 percent. Most of us spend quite a bit of time with the browser open, so just a few little tips can make a big difference in your overall productivity through the course of the day. In this column, I'll share 10 of my favorite IE 9 tips.



## Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).



Email

**⑩ Pin a website to the taskbar**—If you have a few websites that you frequently visit, pinning them to the taskbar can let you jump directly to them from anywhere on your desktop. You can pin a site to the taskbar by clicking the icon to the left of the web address and dragging it to the Windows taskbar.

**⑨ Display tabs on a separate row**—By default, the IE tabs are situated on the same row as IE 9's address bar, which limits both the number of tabs you can see and the length of the address bar (aka the *One Box*). You can display your tabs on a separate row by right-clicking in the tabs area and selecting the option *Show tabs on a separate row*.

**⑧ Quickly open a new tab**—To open a new tab, you don't have to click that little New Tab box. Instead, you can quickly open a new tab by just double-clicking anywhere in the tab area to the right of the open tabs.

**⑦ Reopen closed tabs and your last session**—One of the most frustrating things is when you accidentally close a tab you were working on. Unintentionally closing a tab is particularly annoying if you got to a page after browsing through several links because you might not know the URL you were most recently on. If you ever do this (and I'm sure you do), you can easily reopen any of your last 10 closed tabs by opening a new tab, then clicking the *Reopen closed tabs* link at the bottom. If your system crashed or was rebooted, you can also restore all the tabs from your old session by using the *Reopen last session* link.

**⑥ Download your files to custom locations**—By default, the files you download with IE 9 all go into your Downloads directory. In Windows 7 or Windows Vista, this location is typically found at C:\Users\<username>\Downloads. However, if you do a lot of downloads, you might want to put them in custom locations by clicking the arrow to the right of Save, then choosing Save As in the download dialog box that appears.

**⑤ Quickly display the menu bar**—Did you even know that IE 9 had a menu? Well, it does. You can quickly toggle IE 9's menu on and off by pressing the Alt key. The menu shows File, Edit, View, Favorites, Tools, and Help.

**④ Take advantage of keyboard shortcuts**—In addition to the Alt key, IE 9 provides a number of other useful keyboard shortcuts: Ctrl + L highlights the address bar (One Box); Ctrl + D adds the current web page to your favorites; Ctrl + J opens the Download Manager;

Alt+Home goes to your home page; and Alt+C displays your favorites, feeds, and history.

③ **Jump to previous pages**—Everyone knows that clicking the back arrow pages you back through your recently viewed web pages. However, you can also jump back to a specific page, which can be handy when a given website traps your browser and doesn't let it page back. To jump to a previous page, either left-click and hold the back arrow or right-click the back arrow, then select the page from the drop-down list.

② **Change your default search provider**—Of course, IE 9 defaults to Bing for your search provider, but if you'd rather have Google as your default search provider, you can add Google by simply entering anything in the address bar (One Box), then clicking the Add button at the bottom of the suggestion box. The Internet Explorer Gallery for IE 9 Add-ons will be displayed. Select Google Search Suggestions, click *Add to Internet Explorer*, then select *Make this my default search provider* in the Add Search Provider dialog box.

① **Show more sites on new tabs**—By default, when you open a new tab you'll see two rows of sites listed. To display more rows, open regedit and navigate to HKEY\_CURRENT\_USER\Software\Microsoft\Internet Explorer\TabbedBrowsing\NewTabPage. Then create a REG\_DWORD entry named NumRows, and set the value to a number from 2 to 5. If you don't like editing the registry, you can perform this task (and many more) by using the IE 9 Tweaker Plus tool, which you can get at [The Windows Club website](#). ■

InstantDoc ID 142560

**Internet Explorer holds 53.83 percent of the browser market, followed by Firefox with 22.87 percent and Chrome at 13.19 percent.**



*Don't just be a cloud builder.  
Be a rainmaker.*

IT thought leaders and over 1 billion end users profit from clouds built on a NetApp storage foundation. To make sure your storage architecture is designed to deliver all the rewards the cloud has to offer, visit [NetApp.com/BuiltOn](http://NetApp.com/BuiltOn).

**Go further, faster®**



*Profitable clouds are built on* **NetApp®**



**Find out more at Booth 1505**

[www.netapp.com/microsoftsolutions](http://www.netapp.com/microsoftsolutions)

©2012 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, and Go further, faster are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

# How **Windows Server 2012** Improves Active Directory Disaster Recovery

The process is simpler and faster, but you still have legwork to do

**T**he day I sat down to write this column, springtime in Texas struck with a vengeance, spawning eleven tornadoes across the Dallas–Fort Worth area. Fortunately, we escaped the tornadoes and hail, and we only got several inches of rain. However, the event also spawned an idea for this month’s column: disaster recovery for Active Directory (AD), and specifically how it’s improved in Windows Server 2012 (formerly Windows Server 8).

How will Server 2012 help AD disaster recovery? [I’ve already written about Server 2012’s “virtualization-safe” AD features](#), but it wasn’t until I was listening to the AD team at this year’s MVP Summit that I understood the positive impact these features will have on the forest recovery process. To understand the improvement, you first must understand disaster recovery in the AD world.

## **Forest Recovery: Rare but Possible**

AD is wonderfully fault-tolerant to physical disruptions. Its distributed architecture has the ability to create updates on any domain controller (DC) and have them replicate to other DCs in the domain or forest. This ensures that if a DC or group of DCs is taken out by a local power failure or an act of nature, the domain or forest will



**Sean Deuby**

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel’s core directory services team. He’s been a directory services MVP since 2004.



**Email**



**Twitter**

continue operating with very little impact to the remaining active user population.

There are situations, however, in which you might have to perform a forest recovery. Problems with the root domain in a root-child forest, mismatched schema updates due to extensive replication problems, severe divergence caused by USN rollback due to performing image restores of virtual DCs—these possibilities are extremely uncommon, but they exist. And because your company probably depends on AD for basic functioning, you must have a plan in place.

What if you should have to perform a forest recovery? First, if you're wondering how you'll know when you must do a forest recovery, I have both a long and a short answer. The long answer comes in the form of a [high-level list of possible forest recovery situations on TechNet](#). The short answer is that Microsoft will tell you, because you'll probably have been on the phone with Microsoft Customer Support Services (CSS) for many hours!

## Forest Recovery Steps

At a high level, the procedure (as currently documented in “[Planning for Active Directory Forest Recovery](#)”) involves several steps. Once you've taken all the DCs for the current, failed forest off the network, there are a number of pre-recovery steps you need to perform to prepare the environment.

The next step is to restore one DC per domain using the last known good set of backups from each domain. You're taking backups of at least two DCs in every domain, aren't you? When each DC has been restored (starting with the root, if you have one), connect it to the network. You'll now have a seed forest, with one DC for every domain.

Once created, you need to build out the seed forest as quickly as possible, performing fresh promotions of AD on the existing DCs. I recommend doing a `Dcpromo /forceremoval` versus a complete OS reinstallation to prep the DCs for a fresh `Dcpromo`. Though not well known, `/forceremoval` basically rips out the AD role from the DC

while leaving the OS untouched, so it's far faster when time is of the essence. This buildout phase is by far the most time-consuming part of the forest recovery process, and thus is the place to focus on streamlining. As I said previously, this is a highly simplified process. It ignores little practical considerations, such as the fact that every employee on the corporate network will be hammering these few DCs if you don't take precautions!

For Windows Server 2003, your streamlining options are limited to the procedural and Dcpromo's *Install from Media*. Procedurally, you can do a lot to ensure a speedy forest recovery, and though not technical, solid procedures are extremely important in a situation like this. Remember, this is an all-hands-on-deck situation; everyone's either running around with their hair on fire or drumming their fingers waiting for the hair-on-fire individuals to get AD back up. There's no time to be reading and evaluating TechNet articles and ActiveDir community forums on best practices. You must have solid and tested procedures that both your central AD team and remote operations can follow to the letter so that the restoration proceeds in a well-thought-out manner despite the stress of the moment.

---

**You must have solid and tested procedures that both your central AD team and remote operations can follow to the letter.**

---

## Speeding Up the Recovery Process

With procedures behind us, let's talk about what you can do with Dcpromo. One of the primary actions of the Dcpromo process is the creation and population of the DC's local directory service database. There are several ways you accomplish this depending on what version of Windows Server your DC-to-be is running. The first way that all versions of Windows support is replicating AD objects in from other DCs in the domain. This method might be just peachy for normal operations, but at a time like this you definitely don't want to depend on network connectivity, reliability, and probable congestion to get your authentication infrastructure back up!

What other promotion tricks can we do? All versions of Windows Server after Windows 2000 support Dcpromo's *Install from Media*

(the /IFM option), which promotes a new DC using a system state backup as the source to populate the local directory service database. The advantages of Dcpromo /IFM are that it doesn't require a network, and it's very fast. This is especially impressive for very large databases; at Intel, using IFM cut a 19-hour over-the-network Dcpromo process down to 10 minutes. The requirements to make this work in a disaster scenario are that you must keep a running set of several versions of the system state backup stored on a non-system partition. Further, you must do this on every DC if you don't want to be dependent on the network for copying system state backups around.

Virtualization gives you more options for a speedy forest recovery for current versions of Windows Server—if you're careful. You can't restore a virtualized AD DC from snapshots or image-based backups (i.e., external backups of the VM's hard disks), or bad things might happen, as the TechNet article "[Running Domain Controllers in Hyper-V](#)" explains. My general rule when working with AD and virtualization is: "Don't do anything to AD that it wouldn't expect in a physical environment." You can still use virtualization advantages to speed forest deployment, however. For example, let's say you have a hub-and-spoke network configuration. You could create a generic virtual machine (VM) image in the network hub, with a known good IFM backup loaded on it, then clone that image several times (before you've made it a DC, and thus avoid any AD virtualization problems). Then, perform Dcpromo /IFMs on the cloned images. This will quickly give you a number of DCs in the hub site to support the network load, and branch offices can temporarily authenticate over the WAN until you can rebuild the branch office DCs.

## Windows Server 2012's Boost to Forest Recovery

With that background in mind, how exactly can Server 2012 make AD disaster recovery a much easier process? It centers on the ability to clone Server 2012 virtual DCs. Server 2012 Hyper-V (and soon



VMware vSphere) passes a value—the VM Gen ID—to VMs to tell them if they’ve been subjected to a virtualization activity such as being restored from a snapshot or image-based backup. Thus warned, the DC can take corrective actions to allow it to continue functioning correctly with the other DCs in the domain and forest.

Let’s take this ability to clone DCs and insert it into the forest recovery process. Now, when you need to scale up the seed forest rapidly, all you need to do is clone the seed DCs. Unlike with Server 2008 R2 and earlier, you don’t need to go through any promotion process, so the AD scale-out process (at least in the same networks as the seed DCs) can be very fast indeed. You could speed up the scale-out even more, or simply make it fast in smaller environments, by using differencing disks as a temporary measure. Once you get the initial virtual disk deployed, the difference in disk size for subsequent virtual DCs is quite small (around 200KB) and thus deploying additional DCs to scale out becomes lightning fast. Remember, the point behind this initial scale-out is to get enough DCs operational to allow your users, resources, and applications to begin authenticating and authorizing again; you must finish the AD build-out in its final configuration across your network to bring things back to normal.

AD forest recovery is probably at the top of the “what keeps AD administrators up at night” list, yet I’ll bet only a small number have a documented plan in place. And few of this number have actively tested it, and even fewer test it on a regular basis. Server 2012 will make the process simpler and faster, but only if you’ve done all the legwork first. ■

InstantDoc ID 142825



**Richmond  
V. Baker**

is a supportability program manager for Windows at Microsoft.

Email



LinkedIn



**Roger  
Osborne**

is a premier field engineer in Public Sector services at Microsoft, where he works with federal and civilian clients in Washington, D.C.

Email



LinkedIn



# The Road Warrior's Laptop Build Guide

Create a fully functional portable lab

**A**re you in the field more than you're in the office, with little to no access to a lab for issue reproduction or product testing? Maybe you're looking for an easy way to demonstrate a feature to your customer, but you don't like the idea of lugging around a separate system. If either of these scenarios sounds like yours, then read on.

The purpose of this article is to demonstrate an alternative build option that will allow for multi-product testing and feature evaluation when multiple physical machines or remote virtual host environments are unavailable. This article (and the associated [Sniper-V TechNet blog](#)) covers building a Windows client platform that you can use for corporate applications and productivity, as well as a Windows Server platform, with Hyper-V enabled, that you can use for testing and customer demos.

This configuration is ideal when you're creating documentation, performing demos, or simply verifying a setting or feature. Customers typically have an issue with vendors placing their devices on the production network, and air card signals can also be a challenge.

If you're comfortable with installing OSs (whether from DVD, through Preboot Execution Environment—PXE—booting, or from a USB thumb drive), following Diskpart directions, and backing up and restoring your system, you have all the expertise you need to set up your machine with this clean, functional, multi-boot configuration—a

configuration that also makes backing up and restoring your lab machine a breeze because you'll be dealing with Virtual Hard Disk (VHD) files.

## The Boot Options

Choose the correct boot option from the following list, based on your desired image.

**Installing from PXE boot.** Perform the PXE boot, and press F12. Select a Windows client/server installation image that doesn't use the Operating System Deployment (OSD) wizard. (You need to get to the default Windows installation screen.)

**Installing from DVD.** Verify that the BIOS is configured to boot from DVD. Then, boot from DVD and allow files to load until you get to the default Windows installation screen.

**Installing from USB.** Verify that the BIOS is configured to boot from USB. Then, boot from USB and allow files to load until you get to the default Windows installation screen.

## Creating VHDs

The first thing you need to do is boot up your laptop using a Windows 7 DVD, PXE Boot, or USB bootable drive. At the Windows installation screen, press Shift + F10 to open a command prompt. You're going to clean the existing hard disk and create both VHDs. To do so, follow these steps:

1. Type *diskpart*, and press Enter.
2. Type *list disk*, and press Enter.
3. Type *select disk #*, and press Enter (where # is the number of the disk on which you want to create your VHD).
4. Type *clean*, and press Enter. (Warning: This step will remove references to any data on the disk.)
5. Type *list disk*, and press Enter. (You should see the disk you selected, with all the disk space free.)
6. Type *create part pri*, and press Enter.



### Joe Quint

is an account technology strategist at Microsoft, where he is a trusted technical advisor to several large telecommunications accounts. He specializes in big data.



Email



### Hollis Williams

is a data center specialist at Microsoft, focused on state and local government customers. He specializes in Microsoft virtualization.



Email



Blog

7. Type *format fs = ntfs quick*, and press Enter.
8. Type *list vol*, and press Enter. (You should see a new volume created and selected, formatted as NTFS.)
9. Type *select volume #* (where # is the number of the volume where you want to create your VHD).
10. Type *assign*, and press Enter.
11. Type *list vol*, and press Enter. (A drive letter should now be assigned to the volume.)
12. Type *create vdisk file = "C:\win7.vhd" maximum = 40000 type = fixed*, and press Enter.
13. Type *create vdisk file = "C:\Srv2K8R2.vhd" maximum = 40000 type = fixed*, and press Enter.

Note that in the final two steps, the drive letter and file path should match those of your configuration. Drive size is optional (this example uses 40GB). We recommend setting both to *fixed* size, rather than *dynamic*, to support BitLocker. In the next section, we'll talk about selecting the proper VHD and installing the OS.

## OS Installation

Now that the two VHDs have been created, you'll select the Windows 7 virtual disk, format it, and install the OS. Once finished, you'll have a working Windows 7 installation within the VHD. So, while you're still within the command prompt window, follow these steps to select the VHD and format the virtual partition:

1. Type *select vdisk file = "C:\win7.vhd"*, and press Enter.
2. Type *attach vdisk*, and press Enter.
3. Type *create part pri*, and press Enter.
4. Type *format fs = ntfs quick*, and press Enter.
5. Type *list vol*, and press Enter. (You should see the VHD volume without a drive letter.)
6. Type *list disk*, and press Enter. (You should see the VHD disk in addition to your physical disk.)

7. Type *exit*, and press Enter to end the Diskpart utility.
8. Type *exit*, and press Enter to return to the Windows installation screen.
9. Follow the onscreen OS installation steps as you normally would, making sure to select the 40GB partition as the destination drive.

A note about Windows Server 2008 R2: After you successfully install Windows 7, and you're ready to proceed with installing Server 2008 R2, you'll need to boot your system with the Server media, then press Shift + F10 at the Windows installation screen. Type *diskpart*, and press Enter. Next, you'll repeat the same steps (above), with the only exception being the first step, which you'll replace with

1. Type *select vdisk file = "C:\Srv2K8R2.vhd"*, and press Enter.

## Hyper-V Role

To run multiple OSs concurrently on a single CPU, a hypervisor is required. The Microsoft Hyper-V role in Server 2008 provides this functionality and is easy to implement. To avoid confusion, this article refers to running Hyper-V as a role inside Server 2008 (or later) as opposed to the standalone product labeled as Hyper-V Server 2008. You manage the Server 2008 Hyper-V role via a GUI or PowerShell. In this article, we discuss configuring the GUI.

It's important to note that Hyper-V works only with hardware-assisted virtualization (which is available in processors that include a virtualization option—specifically processors with Intel Virtualization Technology or AMD Virtualization technology) or when hardware-enforced Data Execution Prevention (DEP) is available and enabled.

Hyper-V is installed through Server Manager as a role. On your newly installed server OS, open Server Manager, right-click the Roles node, and click Add Roles. In the Server Roles dialog box, select Hyper-V and Next. The wizard will prompt you to select a primary NIC adapter. Select your primary NIC; however, note that this selection will change



later. Follow the remainder of the prompts to complete the Hyper-V role installation. At the end of the installation, the OS will need to restart.

Once the machine has restarted, Hyper-V will be accessible via the Hyper-V Virtual Machine Manager under Start, Administrative Tools. From this tool, virtual machines (VMs) can be created, deleted, started, stopped, and accessed via a remote desktop console.

## Virtual Networking

Setting up the virtual network based on a wireless adapter is complex. The following five-step procedure uses the 172.16.1.1/24 subnet (can be substituted) and should be followed in order.

1. Enable the wireless adapter feature in the host Server 2008 OS. In Server Manager under Features Summary, click Add Features. Select the Wireless LAN Service, and follow the installation prompts. Configure the wireless adapter to function with your wireless network.
2. Create a VPN based on the loopback adapter in Hyper-V Network Manager. Open the Hyper-V Manager to open the Virtual Network Manager. Create a virtual network by selecting Internal and then Add. In the New Virtual Network dialog box, name the interface EXT-LOOPBACK and click OK. Configure the EXT-LOOPBACK interface via the Network and Sharing Center. To avoid confusing this adapter with any other adapter, rename the interface from Local Area Connection XX to EXT-LOOPBACK. Set the IP address to a private IP in a class C subnet that isn't used within your organization. For example, use 172.16.1.1 with a subnet of 255.255.255.0. Be sure to leave all the other fields empty! No default gateway and no DNS information should be populated.
3. Add RRAS via the Network Policy and NPAS in Server Manager. Open the Server Manager interface and select the Network Policy and Access Services role. As part of the Add Roles wizard for NPAS, choose to install RRAS with Remote Access Service and

Routing. Configure NPAS in Server Manager by right-clicking Network Policy and Access and choose Custom Configuration. In the RRAS Setup Wizard, select NAT and LAN Routing. Now, start the RRAS service.

4. To configure RRAS, expand Routing and Remote Access in Server Manager. Expand the IPv4 node, and right-click NAT to add a new interface. Choose the wireless connection (or primary interface), and select Public Interface with *Enable NAT on this interface* selected. Click OK, and then right-click NAT again, this time selecting EXT-LOOPBACK. In the NAT properties dialog box, select *Private interface connected to private network* and click OK.
5. Configure the VM to leverage the private network. Open the settings for the desired VM and choose Network Adapter. Select EXT-LOOPBACK as the network, and click OK. Internal to the guest OS, configure the network adapter to use any usable IP address on the private network. In this example, we used 172.16.1.1/24, which provides a usable address range of 172.16.1.2-254 and a default gateway of 172.16.1.1.

## Check Out Our Blog

Are you looking for a way to take a fully functional lab with you when you visit your customers? Are you low on hardware and want to use one machine to create a lab environment or even an environment that you're going to run for an extended period of time? Just follow the instructions outlined in this article, and you'll be on your way to having an environment that can be used to replicate most setups. If you'd like to ask questions about this article, peruse additional information that we didn't have enough space to provide in the article, or discover new information that we have regarding this topic, please visit the [Sniper-V TechNet blog](#). ■

InstantDoc ID 142916

OCT 29 - NOV 1 • BELLAGIO • LAS VEGAS, NV

CLOUD  
CONNECTIONS

WINDOWS  
CONNECTIONS

Microsoft  
Exchange  
CONNECTIONS

SQL Server  
CONNECTIONS

SharePoint  
CONNECTIONS

KEYNOTES



**PAUL THURROTT**  
WINDOWS IT PRO  
Senior Technical  
Analyst



**JEFFREY SNOVER**  
MICROSOFT  
Distinguished  
Engineer and the  
Lead Architect for  
Windows Server



**MARK MINASI**  
MINASI RESEARCH  
AND DEVELOPMENT



**MARY JO FOLEY**  
ALL ABOUT  
MICROSOFT  
Editor

# the JOURNEY CONTINUES

JOIN MICROSOFT & INDUSTRY EXPERTS  
AS THEY **HELP YOU NAVIGATE** THE NEW  
AND EXCITING TECHNOLOGIES & RELEASES



REGISTER TODAY! [www.WinConnections.com](http://www.WinConnections.com) • 800.438.6720 • 203.400.6121

# FAQ

## Answers to Your Questions

**Q:** Is there an easy way to view security information about incoming messages in Microsoft Outlook?

**A:** Microsoft doesn't seem to make it easy in Outlook to control some security aspects for inbound messages or to empower us to make decisions about the content of specific messages. We can force inbound messages to render in plain text; however, this setting is for all messages by default, with exemption options for specific senders. We can view header information for a specific message, but the option to do so is hard to find and not very intuitive.

XIntercept has created a small toolkit for Outlook called PocketKnife Peek. (This tool is in addition to the company's PocketKnife product, which provides better integration of Outlook contacts with Microsoft Word.) PocketKnife Peek lets users view aspects of email messages prior to opening them in Outlook itself. Peek lets you perform four actions on a message:

- View the message in plain text
- View the HTML source
- View the Internet header content
- See attachment list, including lists of files within unprotected .zip files



Mike Danseglio



Jan De Clercq



William Lefkovich



Avril Salter



John Savill



Greg Shields



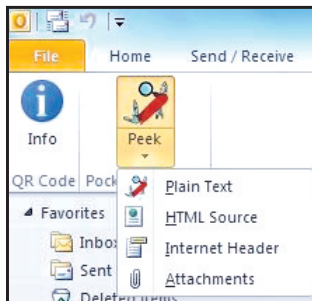
PocketKnife Peek uses a standard installation and can be deployed on Outlook 2000 or later, but only on the 32-bit version of Outlook 2010. As with all Outlook add-ins, Outlook must be restarted for the new add-in to load. After Outlook is restarted, Peek is integrated into the Outlook ribbon in Outlook 2010 and in the toolbar for earlier versions. Peek is also incorporated into the context menu (right-click menu) when an email message is selected. Figure 1 shows the options from the Peek menu in the Outlook 2010 ribbon.

When Peek is initiated on an email message, it opens the message in its own window with a toolbar on top and four tabs on the bottom. The toolbar includes basic email client functionality, such as Reply, Forward, Move to Folder, and Mark as Read. Figure 2 shows a sample message viewed in the Peek interface.

The four security functions of Peek are all easily accessed from the tabs at the bottom of the interface. You can view messages in

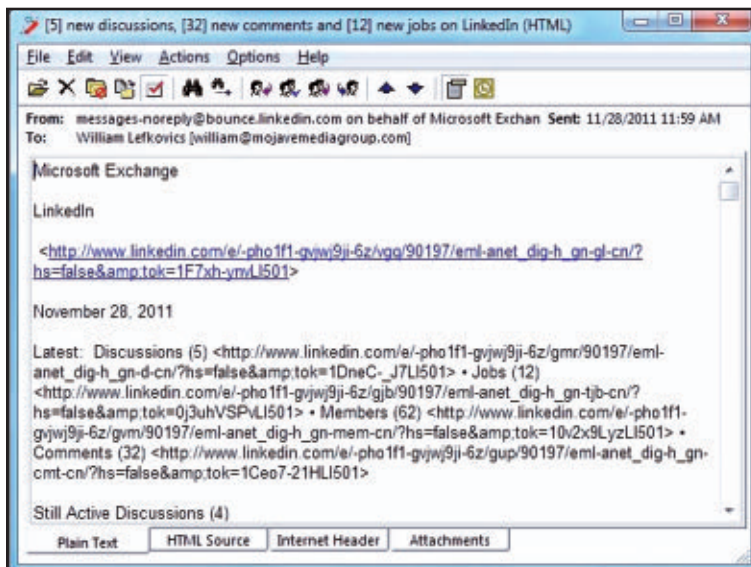
**Figure 1**

Options for PocketKnife Peek in the Outlook 2010 Ribbon



**Figure 2**

A sample message viewed in the PocketKnife Peek interface



plain text, including seeing the actual URL behind any URL displayed in the HTML rendering of the message. Microsoft actually introduced the ability to read all inbound messages as plain text in Outlook 2003 (see the Microsoft article “[How to view all e-mail messages in plain text format](#)”). This feature remains in both Outlook 2007 and Outlook 2010. Peek doesn’t change the original message, but instead lets the user view the message almost on a “what if?” basis. Peek doesn’t mark the email message as Read by default; however, if you use Peek regularly to view message content, you can configure it to mark messages as Read (see Figure 1).

The HTML Source tab lets you read the HTML for the message that has been received. I would like to see external source URLs highlighted in this view or for the interface to use color to distinguish HTML tags from page content, but all the text is black on a white background. The Internet header for the message is easily accessed through Peek as well. The attachment tab lists the attachments to the message, if any, and will also list files in any .zip files that aren’t password protected.

PocketKnife Peek isn’t revolutionary, but it does make viewing message properties in Outlook easier, combining basic security checks into a single interface. Administrators and power users might find themselves using Peek to read all their email.

—William Lefkovich

InstantDoc ID 142695

## **Q:** Which 3 rights are required for the vCenter Server service domain user account?

**A:** VMware vCenter Server 5.0 instance relies on either a Windows domain user account or a built-in SYSTEM account for certain activities. When an Active Directory domain user account is used, that account must be granted three rights on the vCenter Server. It must be a member of the Administrators group, and it must



have the permissions *Act as part of the operating system* and *Log on as a service*. Note that the vCenter Server installer grants the *Log on as a service* right automatically as part of the installation.

—Greg Shields

InstantDoc ID 142547

## **Q:** Can I apply a different password policy to two different Active Directory (AD) organizational units (OUs)?

**A:** No, AD doesn't support different password policies on different OUs—but you can use a workaround that calls on shadow groups, which I'll explain. In Windows Server 2008, Microsoft introduced fine-grained password policies that let administrators apply different password policies to AD user and global security group objects. However, fine-grained password policies can't be applied to an AD OU.

As a workaround, you can use shadow groups to apply a fine-grained password policy to the users that are contained in an OU. A shadow group is a global security group that you “logically map” (meaning that the mapping doesn't require AD configuration changes) to an OU to enforce a fine-grained password policy. To ease administration, you should align shadow group naming with your OU naming scheme.

When using shadow groups, you create a global security group for each OU where you want to apply another password policy, and add the users that are in the OUs as members of the newly created shadow groups. You can then apply different fine-grained password policies to the different shadow groups. Keep in mind that when using shadow groups, if you move a user from one OU to another, you'll also need to update the membership of the corresponding shadow groups.

—Jan De Clercq

InstantDoc ID 142692

**Q:** What tools should I use to detect interference to my Wi-Fi system?

**A:** Wi-Fi interference can be a nightmare, resulting in dropped connections, lowered bandwidth, unpredictable performance, and a host of other problems. Finding it with a spectrum analyzer is incredibly easy, but in most cases not financially practical.

Most IT departments don't need to spend five figures on a true radio frequency (RF) spectrum analyzer. The best bet is to create the device that we dubbed the IT Wi-Fi Tool (ITWiT). We make the ITWiT from surplus components and inexpensive off-the-shelf products. Assuming you're seeing interference in the 2.4GHz band (channels 1-14), the parts list is:

- A surplus laptop running Windows XP or later
- A mini-PCI or USB Wi-Fi card that accepts external antennas
- One or two directional 2.4GHz antennas (optional but useful)
- A USB spectrum analyzer, such as Wi-Spy
- A USB GPS receiver

The ITWiT makes finding interference sources a breeze. You simply fire up the device, start the software, and follow the software indicators until you're literally standing next to the interference source.

—Avril Salter, Mike Danseglio

InstantDoc ID 142425

**Q:** I want to emulate some network devices to use with System Center Operations Manager 2012 network monitoring capability—what can I use?

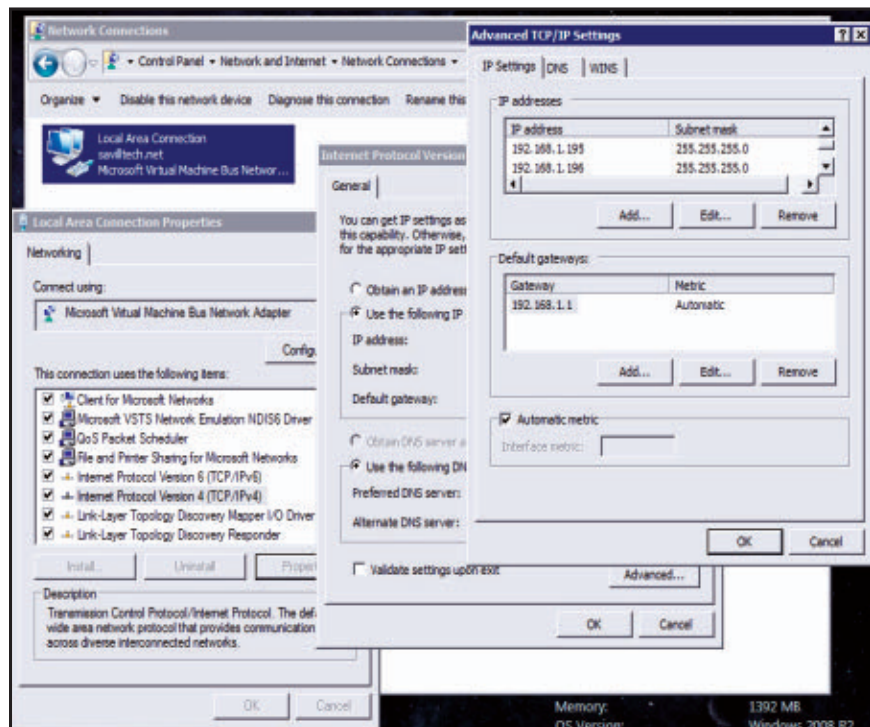
**A:** System Center Operations Manager 2012 offers several new features. Two of the major ones are Application Performance Monitoring (formally known as AviCode) and network device discovery and monitoring.

To test the network device monitoring feature, I originally scouted around eBay and bought a Cisco 1600 series router for around \$30. Then I recently found a free network device simulator at [Jalasoft](#), which, once running, allows several devices to be simulated. They can then be discovered and monitored with Operations Manager 2012.

After the network simulator is downloaded and installed, there are a few steps to get it running. First, don't run it on the Operations Manager server; run it on a spare virtual machine (VM). Next, for every simulated network device, add a new IP address to the network adapter on the OS.

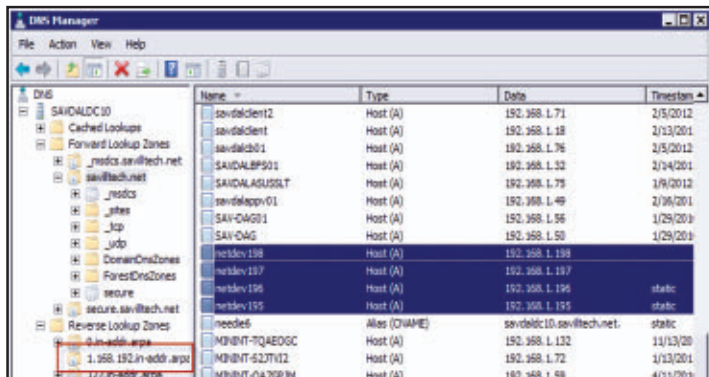
I have four simulated network devices, so I added four additional IP addresses. Open Network Connections, select your network adapter, select Properties, then select Internet Protocol Version 4. Select Advanced, and under IP addresses, click Add; enter the additional IP addresses, then click OK to all dialog boxes (see Figure 3).

**Figure 3**  
Adding additional IP  
addresses



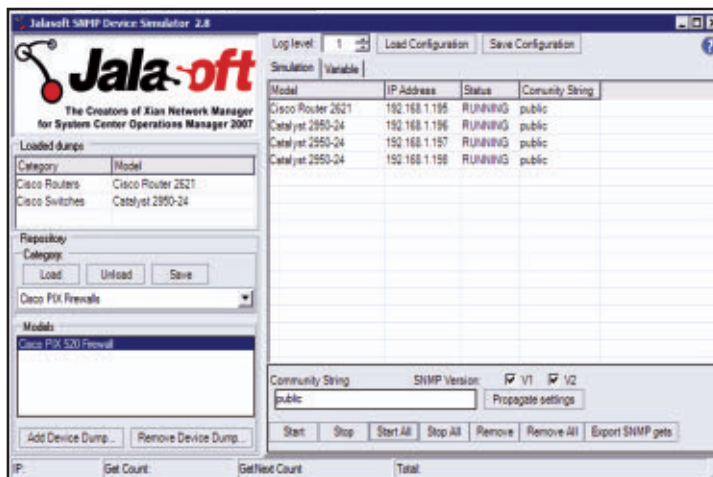
Next, you need to add a host record and reverse PTR record for each IP address with a unique name (this seems to help Operations Manager cope with multiple devices on the same MAC). This is done in DNS, and you need a reverse Lookup Zone for the reverse PTR to be created in (see Figure 4).

Now the simulator can be started, and you can add multiple instances of devices, each using one of the new IP addresses you added. Select a Category from the drop-down menu, then select a Model, and select Load. Right-click the device in the *Loaded dumps* section, select Simulate Device, and select the IP address. The devices can then be started using the Start or Start All action (see Figure 5).



**Figure 4**

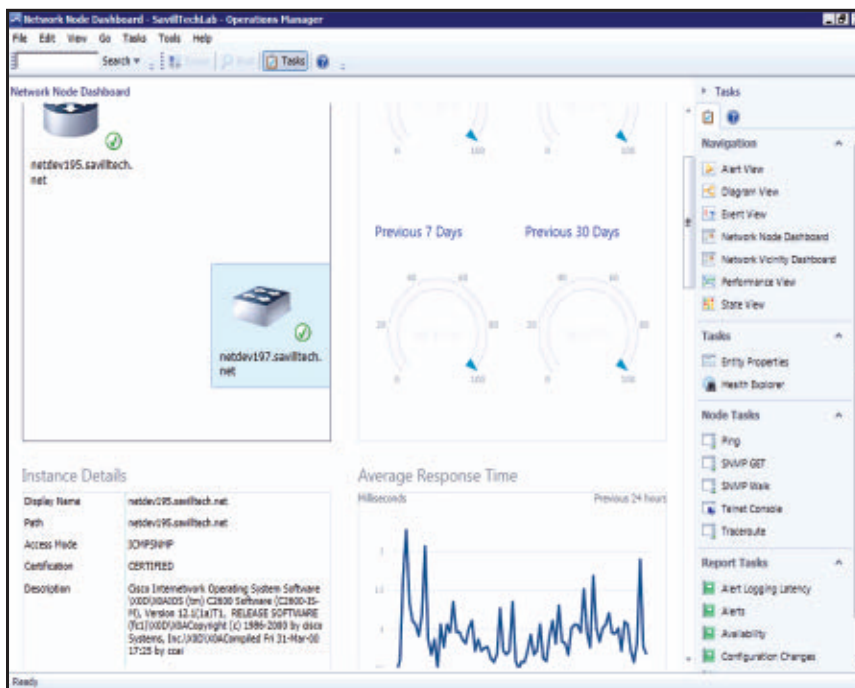
Adding a host record and reverse PTR record



**Figure 5**

Using the Start or Start All actions

**Figure 6**  
Network Node  
Dashboard view



Within Operations Manager 2012, create a new Discovery Rule under the Administration workspace, Network Management, Discovery Rules, and configure it for the correct IP addresses and community string (public). After a period of time, the new devices should show, and their health reporting should display under Monitoring, Network Monitoring, Network Devices. The Network Node Dashboard gives a nice view, which Figure 6 shows. I've observed that sometimes a device doesn't display, but the majority typically do. ■

—John Savill  
InstantDoc ID 142868



# Getting Started with Windows PowerShell

Don't be afraid to jump in

**Y**ou're doubtless aware that Windows PowerShell exists, and that it's somehow a big deal in the Microsoft world. But perhaps you've been too busy to really dig into PowerShell, or maybe you aren't sure why you should care about it. Let's clear up some misconceptions and answer that question in the process.

## Why PowerShell?

First, we need to bust two major PowerShell myths.

**PowerShell is a scripting language.** Not true. PowerShell *contains* a scripting language: a very tiny one, with only about two dozen keywords. But PowerShell is actually a command-line *shell*, much like `cmd.exe` or UNIX Bash shell. You run commands—such as `Ipconfig`, `Ping`, and other commands that you've undoubtedly run in the past—in this shell. True, at some point you might want to combine multiple commands into a batch file, and you're welcome to call that scripting



## Don Jones

has more than 20 years of IT experience, is the author of more than 35 books, and is a speaker at technology conferences such as Microsoft TechEd and Windows Connections.



Email



if you want to. But it isn't really programming such as you'd perform in Microsoft Visual Studio.

Part of this myth comes from the fact that PowerShell allows itself to be used as a kind of lightweight programming language by folks that have those skills. But it doesn't *require* you to have those skills to be effective.

***Microsoft wants to push people away from the GUI.*** Also not true. Microsoft is trying to de-emphasize running GUI consoles on the server; servers aren't well-equipped to deliver a good GUI experience without compromising server-level performance. But running GUIs on the client, even one that connects to a server to do its work, is very much alive.

PowerShell merely represents an alternative for administrators. GUIs can be built to run PowerShell commands in the background, so PowerShell can both provide the functional guts to a GUI and be available for direct use by administrators.

That said, Microsoft is definitely engaging in triage over what its GUI consoles can do, typically focusing on day-to-day tasks. Unusual or less frequent tasks might eventually be available only via the command line. Be upset about that if you want, but also understand that failing to add PowerShell to your skills set could damage your career in the long term.

## Running Commands

Using PowerShell doesn't need to be complicated. For example, suppose you want to add a new user to Active Directory (AD). That task is fairly easy:

```
New-ADUser -Name DonJ -samAccountName DonJ -Title CTO -City "Las Vegas" -Department IT
```

As you can see, the New-ADUser command accepts several command-line parameters. Those parameters (e.g., -Name, -Title, -City) correspond to the same fields that you'd see when using Active Directory

Users and Computers to add a user. So why use the command line instead of the GUI? Because the command line makes it easier to do many things in bulk.

Suppose you've been given a Microsoft Excel spreadsheet of new users that need accounts. The first row of the spreadsheet contains column headers: City, Title, Department, Name, and samAccountName—the attributes of those users. Simply save the Excel file as a comma-separated value (CSV) file, perhaps naming it NewUsers.csv. From there, you can easily use PowerShell to create the new users:

```
Import-CSV NewUsers.csv | New-ADUser
```

As you can imagine, this approach is much faster than manually creating the users in the GUI. PowerShell takes just a few seconds to create 100 user accounts, whereas creating that many accounts in the GUI can take a few hours. By the way, New-ADUser is a command from the Microsoft ActiveDirectory module, which you'll find in the Remote Server Administration Tools for Windows 7 and on Windows Server 2008 R2 (and later) domain controllers (DCs). You need to run

```
Import-Module ActiveDirectory
```

to load the module into memory after installing it on your system.

## Learning the Syntax

The most difficult thing about any command-line interface (CLI) is learning the command syntax. Which parameters are available? What does each one do, and which values will it accept? Administrators often spend hours on a search engine looking for syntax examples. But with PowerShell, you can get started much more easily. Need to do something with services?

```
Help *service*
```

Run that command in the shell and you'll get a listing of commands that deal with services. Suppose you find a promising-looking one in Set-Service. You can quickly learn how to use that command by running another command:

```
Help Set-Service -full
```

The `-full` parameter is important. This parameter provides a lot of extra information, including details on each and every parameter's usage as well as practical examples. No need to hop on Bing to find examples; they're already in the product.

You can use the formal `Get-Help` command rather than just `Help`. I prefer the latter, which automatically pauses the display after each screen of text. There's no need to pipe the file to `More` when you use `Help`.

## A Few Gotchas

PowerShell includes a few gotchas that can trip up newcomers. For example, you need to know that all PowerShell commands run in something called a pipeline. Basically, each time you press Enter, all the commands you just typed are put into a pipeline and run. Most commands accept input and produce output that way. Here's an example:

```
Get-Process | Sort -Property VM -Descending | Select -First 10
```

Whatever the last command in that pipeline outputs is what appears on the screen. If the last command outputs nothing, you get nothing on the screen. Try this instead:

```
Get-Process | Sort -Property VM -Descending | Select  
-First 10 | Export-CSV procs.csv
```

Because `Export-CSV` doesn't produce output, nothing appears on the screen. You do, however, get a CSV file on disk, which is cool.

Also, be aware that when something does appear on the screen, the output's appearance is governed by PowerShell's formatting defaults. You can manipulate those:

```
Get-Process | Sort -Property VM -Descending | Select -First 10 |
  Format-Table -Property ID,VM,PM,Name -autoSize
```

However, the four Format cmdlets—Format-List, Format-Table, Format-Wide, and Format-Custom—don't produce traditional output. They produce instructions that are designed to construct an onscreen display. Therefore, a command such as the following won't work as you might expect:

```
Get-Process | Sort -Property VM -Descending | Select -First 10 |
  Format-Table -Property ID,VM,PM,Name -autoSize |
  ConvertTo-HTML | Out-File procs.html
```

The input to ConvertTo-HTML is screen-layout instructions, so you'll get an HTML file with those instructions. It's pretty much just all hexadecimal codes and garbage, as far as we human beings are concerned. There's an easy way to avoid this gotcha: Don't put any command after a Format cmdlet, unless that command is Out-Printer or Out-File, both of which are specially designed to understand screen-layout instructions.

## Extending the Shell

Like the Microsoft Management Console (MMC), PowerShell is designed to be extended so that it can manage different technologies. As Table 1 shows, there are two ways to extend the shell, depending on whether you're using PowerShell version 1 (v1) or version 2 (v2). Both methods offer a means of discovering which extensions you have installed locally.

Note that the PowerShell v2 method of finding what's available lists only modules that are installed in specific locations; some product

Table 1: Extending PowerShell

Task	PowerShell v2 Only	PowerShell v1 and v2
Find what's available	Get-Module -ListAvailable	Get-PSSnapin -Registered
Load an extension	Import-Module <i>name</i>	Add-PSSnapin <i>name</i>
See which commands were added	Get-Command -Module <i>name</i>	Get-Command -PSSnapin <i>name</i>
Learn to use a command	Help <i>command-name</i>	Help <i>command-name</i>

extensions don't put their modules in the right place. However, those products typically create a Start menu shortcut that will tell you where the module is installed.

Some folks get a bit hung up on the idea that there are product-specific versions of PowerShell, such as the Exchange Management Shell or the SharePoint Management Shell. The fact is that there's no such thing as product-specific versions of PowerShell, although Microsoft named features in a way that implies otherwise. The Exchange Management Shell is nothing more than a Start menu shortcut. Look at its properties: You'll see that it runs good old PowerShell.exe. The shortcut simply simultaneously runs PowerShell and auto-runs a particular script, or loads a particular module. You can manually load that module: Just look at the properties of the Start menu shortcut to see the module location, and then run

```
Get-Module path-to-module
```

That way, you can load whichever modules you like into the shell.

## Objects

There's a lot of talk about objects in PowerShell, which freaks out some people. They immediately think, "This is starting to sound like development, and I didn't sign up for programming. Count me out!"

Chill. “Object” is just a word that means “a data structure.” Imagine an Excel spreadsheet or even a Microsoft Access database table. Each row in the table or sheet is an object; each column is a property. It’s that simple. The command

```
Get-Service
```

produces an onscreen table display, with each row representing one service object, and each column representing an object property. This isn’t programming; it’s just a different terminology. With that terminology in mind, you can do some pretty amazing things.

***Remove objects you don’t want to look at.*** Use the Where-Object cmdlet to filter things out of the pipeline. Here’s an example:

```
Get-Service | Where { $_.Status -eq 'Running' }
```

This command displays only running services. Within the curly brackets (which specify the criteria for objects that you want to see), you use `$_` to refer to whatever the previous command produced. I wanted to look at only a portion of the entire service object. How do you specify a portion, or fraction, in math? With a decimal point, as in 3.147, right? So I followed `$_` with a decimal point, and then specified the fraction with which I wanted to work: Status. I know from previous runs of `Get-Service` that the Status column contains “Running” for running services, so I said I only wanted to keep the services in which the Status column equals (that’s the `-eq` comparison operator) “Running.”

***Choose which properties to display.*** You can easily choose columns other than the defaults:

```
Get-Service | Select -Property Name,Status
```

***Change the sort order.*** The default sort order is always ascending. You can reverse that if you want to:



```
Get-Service | Sort -Property Name -Descending
```

**Combine.** You can pipe the output of many commands to the input of another:

```
Get-Service | Sort -Property Status |  
Select -Property Name,Status
```

## Don't Let the Internet Stress You Out

One problem with the Internet is that it gives *everyone* a voice. PowerShell is designed to be approachable and usable by many audiences. But just because you've found some 900-line sample script doesn't mean that kind of programmer-style approach is the only way to use PowerShell; it's just how one person chooses to use it, probably because that's the approach that person understands the best.

On <http://ShellHub.com>, I've collected a short list of PowerShell resources that tend to favor a more administrator, type-a-command-and-press-Enter approach, rather than the programmer-style approach favored by some enthusiasts. When you're more comfortable with the shell, you'll likely start modifying your own approach, getting ever-more complex. That's fine. But you can start simple and still do great stuff. For example, in my book *Windows PowerShell Scripting and Toolmaking*, I start with a simple command that you can type in to get immediate results. I gradually build on that command, adding parameters, adding help, and so forth, until 100 pages later it looks and feels like a native PowerShell cmdlet. You don't need to tackle all PowerShell's complexity at once. Start small and work your way up.

## You Don't Need to Start Over

I'll often ask a trick question of my students when I teach a PowerShell class: "How do you map a network drive in PowerShell?" Many will search through Help, struggling to find something. They'll often run across New-PSDrive, which isn't the right answer; its drives aren't

visible outside of PowerShell. Eventually I'll give them the answer: Net Use.

Everyone smacks their foreheads. It's a good lesson: Microsoft isn't asking you to discard everything you already know. (Weird, right?!) All the command-line techniques you already know—Net Use, Icacls, DsAcls, NSLookup, Ping, Ipconfig, Pathping—still work, so continue to use them. You can mix and match those pretty freely with PowerShell's native cmdlets, too. So if you already know how to do something, don't bang your head against a wall trying to figure out the PowerShell way to do it. Just do what you know.

## Why Is PowerShell Scary to Some Admins?

I'll tell you a little secret. Many administrators got into Windows administration, as opposed to UNIX or something else, specifically because Windows is *easy*, at least on the surface. Run through a few wizards, click a few buttons, and your job is done. Many of those administrators (I'm sure you've met a few) have very little understanding of what's going on underneath those buttons and wizards. That's why PowerShell frightens them. It isn't that they're afraid of learning the syntax, or even that they think typing is so tedious. It's that PowerShell strips away much of the hand-holding that the GUI has been doing.

Being able to use PowerShell well requires a firm understanding of what's happening in the technology you're administering. PowerShell expertise is a symptom of an excellent, knowledgeable administrator. Unfortunately, years of having the details hidden by a GUI has left many of our colleagues less knowledgeable than they ought to be. I've spoken with administrators who've never run Ping, who couldn't troubleshoot AD replication without a GUI tool, and who know little about how email messages are queued and delivered on Exchange. Those admins have a tough time using PowerShell.

Here's a question for you: Can you, or can you not, create an AD user account that has a blank logon name? The answer is, *you can*—once. The user won't be able to log on, of course, but AD's only requirement

for a logon name is that it be unique. “Blank” is unique the first time you use it; it’s only when you try to create a second blank user that the operation will fail. Of course, if you’ve interacted with AD only through the GUI console, you won’t know this because the GUI console won’t let you create a blank logon name at all. But if you really know what AD is doing under the hood, then the answer is obvious.

PowerShell demands this kind of in-depth skill and knowledge. The good news is that having that in-depth skill and knowledge also makes things such as troubleshooting, architecture, planning, and so on, much easier. Getting yourself to the point where you can use PowerShell effectively *will* make you a better administrator, not because of PowerShell, but because of the greater technical knowledge it requires. So dig in.

### Use It or Lose It

I have a saying that’s become something of a proverb: “Learn PowerShell or learn to say, ‘Would you like fries with that?’” I’m already seeing organizations discriminate on PowerShell. If PowerShell expertise is a sign of a better, more knowledgeable admin, why not choose to keep and promote those folks over less knowledgeable people? PowerShell expertise is a more reliable indicator of true expertise than any certification exam that Microsoft could create. If you think certifications were a key to you getting and maintaining your job, then make no mistake: PowerShell is even more crucial. Ignore it at your peril.

Here’s another way to think about it: If your sole value to your company is your ability to click “Next, Next, Finish” without understanding what’s happening under the hood and without being able to automate that task to make it faster and more efficient . . . well, you’re what I call a “button monkey.” In other words, you’re easily replaced. Me, I’d rather be the one with the “arcane” knowledge of how to automate things so that I don’t have to spend my time doing boring, repetitive tasks—and so that I’m much harder to get rid of. ■

InstantDoc ID 142467

# Windows Server 2012 Storage Spaces

This feature gives you a new way to administer storage

**T**he core file system capabilities in Windows Server have not changed radically since the earliest version of Windows NT. Windows 2000 introduced dynamic disks. There have been improvements to NTFS reliability and performance. File services have had incremental changes, as has the file-sharing Server Message block (SMB) protocol, but nothing groundbreaking. File Server Resource Manager (FSRM) added capabilities around screening, quotas, reporting, and classification but didn't change the core capabilities of the file system and how file services are used.

When I think of file services, I think of a server on which to store your Microsoft Word and PowerPoint documents. If this server needs to be highly available, it can be clustered, with one server at a time offering the file share. Volumes with fault tolerance can be created on dynamic disk by using Windows mirroring (RAID 1) or striping with parity (RAID 5) capabilities. However, IT administrators must manually perform both the selection of disks and any repair actions. Furthermore, more advanced features, such as thin provisioning of storage and the easy addition of more storage to a pool in which volumes can be created, just aren't possible without the use of a separate SAN or NAS. But this changes in the next version of Windows Server, Windows Server 2012 (formerly code-named Windows Server 8).

## Using Storage Spaces

As an IT administrator or even an end user, you often need storage that might sometimes require fault tolerance. Other times, you just need to be



### John Savill

is a Windows technical specialist, an 11-time MVP, and an MCITP: Enterprise Administrator for Windows Server 2008. He's a senior contributing editor for *Windows IT Pro* and is currently writing his latest book, *Microsoft Virtualization Secrets* (Wiley).



**Email**



**Twitter**



**Website**

---

**Storage Spaces, combined with the in-box iSCSI and SMB 2.2, provides a powerful storage solution for servers and desktops alike.**

---

able to store information that is protected in another way. You can open the Disk Management Microsoft Management Console (MMC) snap-in, examine the physical disks, convert them to dynamic disks if necessary, and then create a volume that meets your requirements. If the volume needs to grow, you might be able to extend it (depending on the physical disks), but you can't add additional disks to an existing volume, to provide easy scalability. For small and midsized organizations or even large organizations with smaller remote locations (i.e., locations with just a couple servers and for which neither SAN nor NAS is economical), providing a good storage solution for services is a huge problem. At the other end of the scale, power users on desktops also struggle to organize their data across internal drives and USB-connected disks.

Storage Spaces is a new feature in Server 2012 (and the Windows 8 client). This feature enables a new way to think about and administer storage. With Storage Spaces, the physical disks that provide underlying data storage are abstracted from the process of requesting new volumes, now known as spaces. The Storage Spaces technology automatically performs any necessary actions to restore data redundancy if a disk fails, provided that sufficient physical disks are available.

The first step is to create a storage pool, which is a selection of one or more physical disks that are then pooled together and can be used by the Storage Spaces technology. Storage pools support USB, Serial ATA (SATA), and Serial Attached SCSI (SAS) connected disks in a Just a Bunch of Disks (JBOD) scenario. With no hardware-based high-availability support such as RAID happening behind the scenes, Storage Spaces takes care of fault tolerance. The use of USB-connected drives is great on the desktop; servers focus on SATA- and SAS-connected drives. In addition, Storage Spaces fully supports shared SAS. You can connect a disk enclosure to several hosts in a cluster, and the Storage Space on those shared SAS drives will be available to all nodes in the cluster and can be used as part of Cluster Shared Volumes (CSV). If you use an external disk enclosure, then Storage Spaces supports the SCSI Enclosure Services (SES) protocol, which enables failure indications

on the external storage. For example, you could enable a bad disk alert if Storage Spaces detects a problem with a physical disk.

Other technologies, such as Microsoft BitLocker Drive Encryption, can also be used with Storage Spaces. When a new storage pool is created, the disks that are added to the storage pool disappear from the Disk Management MMC; the disks are now virtualized and used exclusively by Storage Space technology. You can see the disk state in the storage pools view in File and Storage Services in Server Manager (on a Windows 8 server ) or by using the Storage Spaces Control Panel applet (on a Windows 8 client). This article focuses on using Storage Spaces on the server with Server Manager and Windows PowerShell, but all the features that I write about here are also available on the client. (The only difference is that on clients, the Storage Spaces applet, instead of Server Manager, is used for management.)

Start Server Manager. Make sure that the server you want to manage has been added to the list of servers on your Server Manager instance (see the sidebar “Windows Server 2012 Management”), then open File and Storage Services. Select the target server from the Servers tab, and then select the Storage Pools tab, which shows information about existing storage pools and disks that can be used in a storage pool. (These disks are system disks that aren’t hosting any volumes.) These unused disks are shown in a Primordial Storage Space and are the building blocks from which storage pools and Storage Spaces can be created, as Figure 1 shows. To create a storage pool, follow these steps:


1. From the Tasks menu, select New Storage Pool to launch the New Storage Pool Wizard.
2. Enter a name and an optional description for the new storage pool, and then click Next.
3. On the next screen, select the physical disks that are available to add to the new pool. Also select the disks’ allocation (Data Store, by default). You can allocate the disks as part of the virtual disks that you will create later or reserve them as hot spares, as Figure 2 shows. Click Next.



## Windows Server 2012 Management

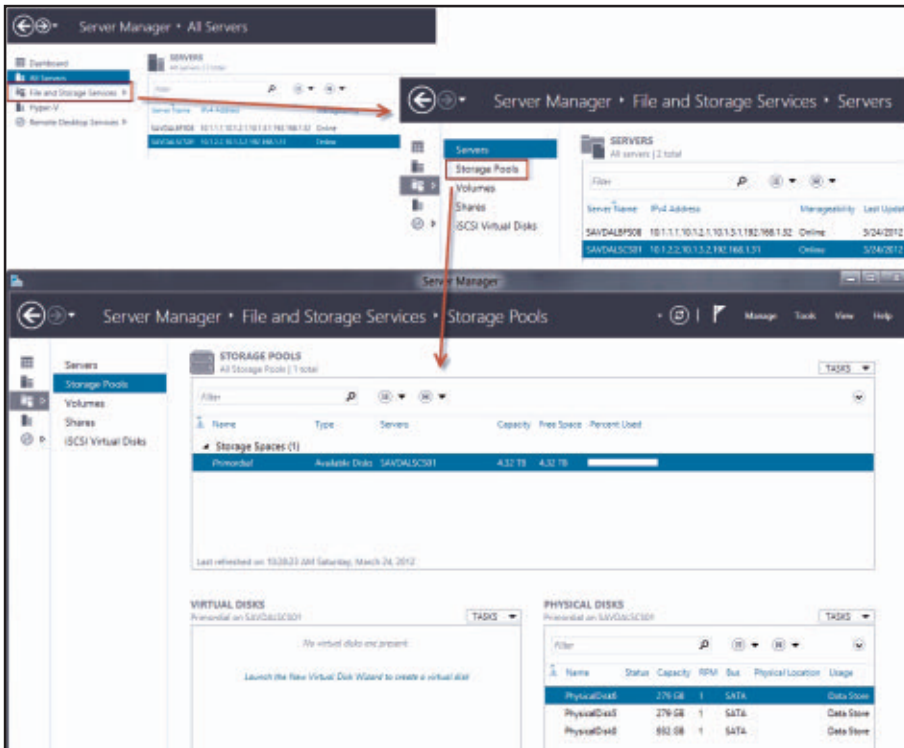
The next version of Windows Server, Windows Server 2012, embraces the management philosophy “the power of many, the simplicity of one.” Even with virtualization, most server management is still performed either by connecting to the OS for Remote Desktop management or by remotely connecting to one server at a time via Server Manager. With Server 2012, all management can and should be performed remotely. After installing the Remote Server Administration Tools for Server 2012 on a Windows 8 client machine, start Server Manager, and create groups of servers, which you can then manage as an entire group. Dashboard views (such as the one in Figure A) make it easy to see any problems on any server in the group and to perform actions on multiple servers simultaneously. Typical management actions, such as adding or removing roles and features, are possible, as are configurations such as Storage Space actions. This doesn’t mean that you can’t run Server Manager locally on servers. But as the trend to optimize management continues, managing a single server makes less sense. As more servers run Server Core instead of a full installation, management will need to be performed remotely unless you want to use Windows PowerShell for all administration.

**Figure A**  
Server Manager

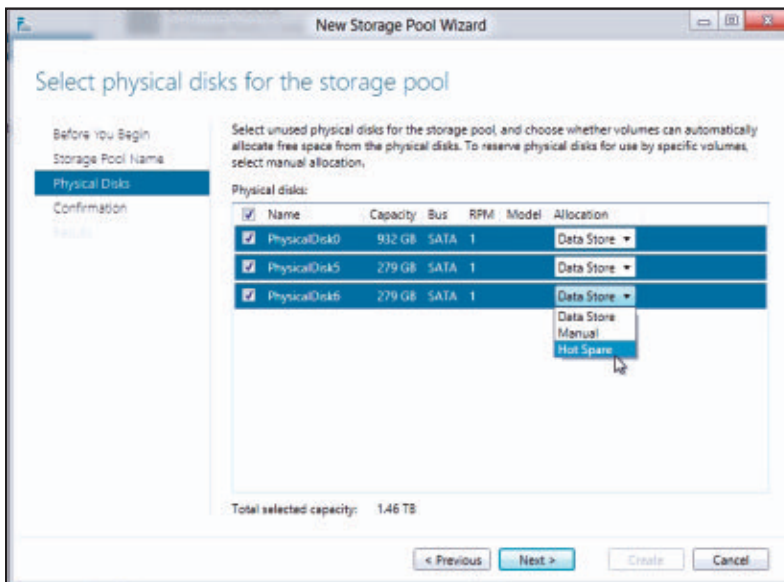


4. Read the displayed confirmation. Click Create to complete the storage pool creation.

A storage pool is now available. The next step is to create virtual disks within the storage pool. You can then create volumes on those disks so that the OS can use them.

**Figure 1**

Remote management through the Server Manager tool

**Figure 2**

Adding and allocating disks to a new storage pool

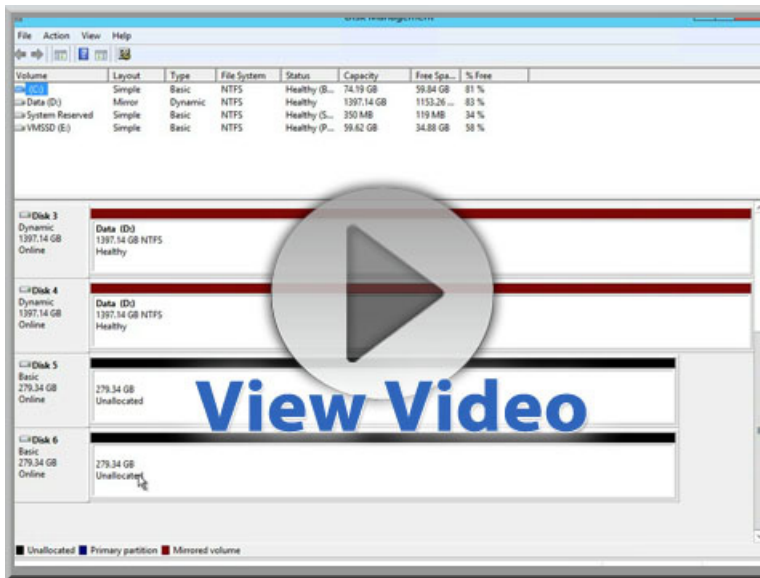
Storage Spaces introduces a feature that was previously available only when using external storage solutions such as SANs and NAS devices: the ability to thin-provision storage. During the creation of a virtual disk, you have two options. The first is to create the disk as fixed, meaning that all the space for the size of the virtual disk is allocated during its creation. The second is to create the disk as thin, meaning that space is taken from the pool only as needed. Using a thin-provisioned disk, you can create a virtual disk that is much larger than your actual available storage. Now, this capability doesn't mean that you can store more data in the thinly provisioned disk than is actually allocated to the pool. But volumes typically fill up over time. I might create a 10TB thin disk that initially has only 1TB of associated physical storage; as the amount of data increases and approaches 1TB, I can add another 1TB of physical storage to the pool simply by adding more disks. As the data approaches 2TB, I can add another 1TB of storage by adding still more disks, and so on. As long as I add physical disks before the virtual disk fills, there's no issue. Alerts can be generated to notify me that a storage pool is reaching its threshold, giving me time to add the required storage.

When you create a virtual disk, all you need to know is in which storage pool to create the disk. No knowledge of physical disks is required or even openly available. The point of Storage Spaces is to create virtual disks as needed. To create a virtual disk, follow these steps:

1. Select a storage pool in which to create a new virtual disk. In the Virtual Disks section, select the New Virtual Disk task.
2. Confirm that the correct server and storage pool are selected in the storage pool selection page of the wizard, and then click Next.
3. Give a name and optional description for the new virtual disk, and then click Next.
4. Select the storage layout, which can be simple (i.e., no data redundancy and data striped over many disks), mirrored (i.e., data duplicated to additional disks), or parity (i.e., spread data over multiple disks but add parity data to help protect against data loss during a

- disk failure). Prior to Storage Spaces, these layouts would have been referred to as RAID 0, RAID 1, and RAID 5. That nomenclature isn't used with Storage Spaces layouts because of differences in implementation. Make your selection, and then click Next.
5. Select the provisioning type (i.e., Thin or Fixed) and then click Next.
  6. Specify a disk size. If you choose Thin as the provisioning type, you can select a larger size than the available physical free space. Click Next.
  7. A confirmation is displayed; click Create.

After the virtual disk is created, it is available in Server Manager and the Disk Management MMC, where you can create volumes and format the disk with a file system. You can see the actual amount of space that the virtual disk uses in a storage pool in Server Manager (or in the Storage Spaces Control Panel applet on a client). See the accompanying video for a walk-through of this process.



## Video

Creating and using  
Storage Spaces

You can also use PowerShell to manage Storage Spaces. For example, to create a new storage pool that uses three physical disks, I can use the following commands:

```
$phyDisks = Get-PhysicalDisk
$storSub = Get-StorageSubSystem
New-StoragePool -FriendlyName "Stuff" -PhysicalDisks
    $phyDisks[0] , -
    $phyDisks[1], $phyDisks[2] -StorageSubSystemFriendlyName
    $storSub.FriendlyName
```

To create virtual disks in the pool, I can use these commands:

```
New-VirtualDisk -StoragePoolFriendlyName "Stuff"
    -ResiliencySettingName Mirror -Size 10TB -Provisioningtype
    Thin -FriendlyName "Data1"

New-VirtualDisk -StoragePoolFriendlyName "Stuff"
    -ResiliencySettingName Parity -Size 10TB -Provisioningtype
    Thin -FriendlyName "Data2"
```

I can output the results of the Get-VirtualDisk cmdlet as a list to get details about a virtual disk. For example, I can use the following command to get information about the number of data copies for a mirror and its operational status:

```
PS C:\ > Get-VirtualDisk -FriendlyName Data1 | fl
```

Figure 3 shows a small part of the resulting output.

**Figure 3**  
Get-VirtualDisk Cmdlet  
Output

NumberOfAvailableCopies	: 0
NumberOfColumns	: 1
NumberOfDataCopies	: 2
OperationalStatus	: OK

## Server Message Block: Better than Ever

In a future article, I'll go into detail about the new storage protocols in Windows 8. But if I want to explain why I think of Server 2012 as a great file services platform, I need to talk at least briefly about how to use the great new Storage Spaces enabled volumes.

Windows has used SMB as its protocol of choice for remote file access for a long time. Server 2012 introduces SMB 2.2. Although seemingly only a minor version increase from the SMB 2.1 in Server 2008 R2, SMB 2.2 actually makes a vast change to both the performance and capabilities of the SMB protocol.

The overall performance of SMB has been greatly improved, making access to data via SMB essentially equivalent to direct access to the storage. This access is enabled by several changes, including SMB Multi-Channel, which now allows multiple TCP connections to be established over multiple NICs if available for a single SMB session. This change enables bandwidth aggregation because multiple NICs and CPUs can be used for network processing when Receive-Side Scaling (RSS) and multiple NICs are leveraged. This also works for Server 2012 native NIC teaming. (Yes, Server 2012 has an inbox NIC teaming solution!) High availability to file shares with failover clustering has also dramatically improved with a new Active-Active mode, which enables a CSV that's accessible to all nodes in a cluster. Multiple hosts in the cluster can simultaneously share the CSV for key workloads such as Hyper-V virtual machines (VMs) stored on a file share or even Microsoft SQL Server databases. This Active-Active file sharing allows for zero downtime and no loss of handles in the event of a failover.

Beyond SMB, Server 2012 has iSCSI as a role service, as part of File Services in File and Storage Services. After it's installed, this service allows a Server 2012 server to act as an iSCSI target, enabling access to its storage from both a file level (using SMB) and a block level (using iSCSI). iSCSI targets on a server are typically virtual hard disks (VHDs), and full configuration of access and authentication services is possible.



## The Tao of Chkdsk

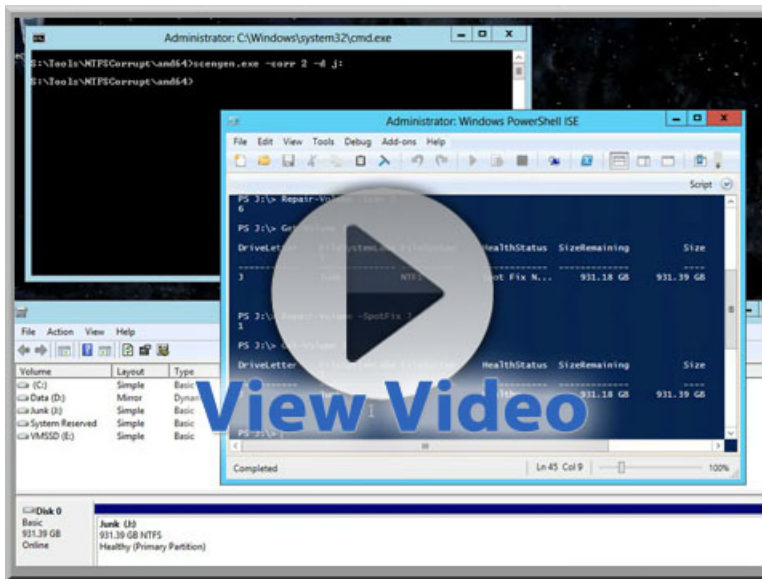
There is still one concern related to the use of Windows Server file services, specifically NTFS, which has very large volumes with many files. When something goes wrong, you might need to run the Chkdsk utility to repair the problem. Chkdsk is very good at its job, but it's a long laborious job, which must go through all the disk content looking for problems, and then perform the repair, which—due to the nature of disks and their speed—can take a very long time (possibly days for large volumes with many files). The result is days during which the volume is offline while the repair operation is performed. Along with considerations of performing a data restore after a disaster, this action is why many times NTFS volumes are kept below a certain size: to ensure that Chkdsk can be run in a reasonable period (i.e., a few hours) if the worst happens.

The new Resilient File System (ReFS) file system, which will become more prevalent in future versions of Windows, aims to reduce the chances of corruption. NTFS itself has become more resilient, with self-healing capabilities, but Chkdsk is still needed at times. Windows 8 has solved once and for all the concerns about running Chkdsk on even the largest volume.

Chkdsk is slow. As I already mentioned, the tool must go through the entire disk and all its content looking for problems, which takes time. As it finds the problems—which will be on a minuscule number of actual files—Chkdsk fixes them. These fix operations take almost no time (i.e., seconds). The problem is that Chkdsk takes the volume offline, making the content unavailable as it performs the health checking and fixing.

Server 2012 breaks the Chkdsk fix process into two parts. The first part scans the disk and data, looking for problems. If a problem is found, then that problem is marked and noted as requiring fixing. The big difference is that the volume is still online, whereas the long search-and-checking process is performed because no fix is actually being performed. Once the scan is completed, if problems need to be fixed, and Chkdsk is run again in a spotfix mode, which takes the volume offline as it performs the fixes on the identified problems that

were found during the scan. The volume is now only offline for seconds instead of hours or days. The scan process has been separated from the actual repair process. Using Chkdsk, the two commands are as follows. The first one will take a long time as it's performing the scan, but there will be no impact to volume availability. The second command would take the volume offline or trigger at next reboot. The accompanying video provides a full example.



## Video

Performing a spotfix on a volume

Chkdsk /scan J:

Chkdsk /spotfix J:

If you're using PowerShell, use the following commands:

Repair-Volume -Scan D

Repair-Volume -SpotFix D

If CSVs are used, there is actually zero downtime when running the spotfix action. The reason is that CSV adds another level of indirection

between the disk and how it's accessed. Also, CSV can actually pause I/O operations for about 20 seconds. This means that when the spot-fix action is run, CSV just pauses I/O to the volume while it's taken offline and fixed. This means as far as users of the CSV volume are concerned, there was just a slight delay in access and no actual off-line action or loss of file handles.

### A Powerful Storage Solution

The Storage Spaces functionality, combined with the in-box iSCSI and SMB 2.2, provides a powerful storage solution for servers and desktops alike. When other technologies, such as the new data de-duplication feature, are also considered, the use cases for the Windows file services platform explode. ■

InstantDoc ID 142786



## Earn your degree and IT certs at the same time! *Online.*

- **Relevant Degrees AND Certifications—**  
Accredited degree programs that incorporate up to 18 certifications at no additional cost.
- **Opportunity to Advance Quickly—**  
An innovative learning approach to education that lets you leverage prior experience and your IT certifications to complete your degree faster.
- **Flexible Online Learning**



[www.WGU.edu/ITPro](http://www.WGU.edu/ITPro) 1.800.264.2995

**WESTERN GOVERNORS UNIVERSITY**

ONLINE. ACCELERATED. AFFORDABLE. ACCREDITED.

# Corrupted Items and Mailbox Moves in Exchange Server 2010

Tweak MRS settings to avoid interruptions

**T**he words “corrupted mailbox data” tend to make Microsoft Exchange administrators break out in a cold sweat. After all, the first thing that those words bring to mind is the possible necessity of a database restore. And we all know how much fun *those* are.

The [TechNet Exchange 2010 Server forum](#) contains multiple discussions about encountering bad or corrupted items during mailbox moves and the resulting fuss when administrators see messages such as *Error: This mailbox exceeded the maximum number of corrupted items that were specified for this move request*. This error appears when administrators run the [Get-MoveRequestStatistics cmdlet](#) to check the status of a mailbox move that’s in progress. The error can also be seen logged as event 1100 from the source MExchange Mailbox Replication in the Application event log.

Just what causes such problematic items? How worried should you be if you encounter this error? And what can you do about the issue? Read on to find out.

## Mailbox Replication Service and Bad Items

Conceptually, the error is easy to understand. The Exchange 2010 Mailbox Replication Service (MRS) runs on Client Access servers and is responsible for moving mailbox contents asynchronously. In other words, MRS moves mailboxes behind the scenes so that users can continue to work



**Tony Redmond**

is a contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press).



**Email**



**Twitter**



**Blog**

---

**MRS can play a role in keeping your Exchange databases in good health.**

---

until their new mailboxes are available and MRS switches the AD points for the user account to point to the new mailbox location. MRS can deal with Exchange Server 2003, Exchange Server 2007, Exchange 2010, and Exchange Online (part of Microsoft Office 365), so it's a pretty capable component. As we'll discuss later, MRS can also play a role in keeping your Exchange databases in good health.

MRS processes mailboxes according to move requests. A move request describes the target database and other parameters that influence how MRS should deal with the mailbox. For example, one of the parameters (`SuspendWhenReadyToComplete`) for a move request can be used to suspend the move after the mailbox contents are nearly fully copied. This action allows an administrator to review the move report, decide that everything is OK, and then resume the move request to allow MRS to flip the Active Directory (AD) settings, and redirect the user to his or her new mailbox before resuming the move. Another parameter indicates the permitted bad item limit—in other words, the number of corrupted items that MRS can tolerate before completing a mailbox move request. In this context, “corrupted” means that MRS determines a property or properties of the item to be incorrect, malformed, or otherwise unacceptable to Exchange.

The default value for the bad item limit is 0, meaning that MRS will tolerate no data loss in a mailbox move. This is sometimes appropriate when data absolutely must be retained. However, you might find that you can't move mailboxes at all because MRS continues to encounter corrupted data and fail with the error message mentioned earlier. For this reason, best practice is to set a reasonable limit for bad items. You should also check mailbox-move reports after requests complete, to determine whether the limit is in fact reasonable or if valuable data is being lost. I think anything in the 10-to-20 range is reasonable, but that depends on your circumstances and organizational requirements.

If you decide that a mailbox really must be moved no matter how much data is lost en route, you can set a very high bad item limit,

which forces MRS to move the mailbox at all costs. The maximum setting is 2147483647, but 100 is probably a better starting point. Note that after you specify a bad item limit equal to or higher than 51, you also need to pass the `AcceptLargeDataLoss` switch parameter to tell MRS that it's acceptable to drop all corrupted items.

## Positive Intervention Is Required

If a move request is halted because MRS encounters too many bad items, an administrator can intervene to amend the move-request parameters and get things moving again. Use the [Set-MoveRequest cmdlet](#) to amend the `BadItemLimit` parameter and increase the tolerance for bad items until MRS can move the mailbox, and then instruct MRS to resume the move.

For example, suppose that the move for TRedmond's mailbox has stalled because MRS has encountered too many bad items. The first order of business is to view the move report. Use the Exchange Management Console (EMC) to view the properties of the move request, find the items that are causing trouble, and perhaps figure out whether you need to intervene in the mailbox before continuing. See the Microsoft article "[View Move Request Properties](#)" for details about how to access a mailbox-move report.

Suppose that the original bad item limit for the move request was 10, and MRS encountered more than 10 corruptions as it copied data from the source to the new target mailbox. We can scan the mailbox-move report to discover how far MRS managed to get before reaching this limit. MRS creates checkpoints on a frequent basis so that it can resume a move request without redoing work; MRS reports this checkpoint status in terms of percentage complete. This percentage reports the overall progress of the move and takes into account other processing, such as connecting to source and target servers, setting up the target mailbox, and preparing to copy. The phase between 25 and 90 percent of the overall move is when copying of mailbox data occurs. If MRS encounters a bad item limit, it will happen at some point in this range.



In our example, suppose that the mailbox-move report indicates that MRS reached 80 percent before stopping the move. The last 10 percent or so of a mailbox move is used to perform a final incremental synchronization of mailbox contents and to switch the mailbox pointers in AD. At 80 percent, there isn't much more mailbox data to copy, so there's a fair chance that MRS can finish copying data from the source mailbox if you increase the bad item count to 20. To do so, you can run the following commands in the Exchange Management Shell (EMS):

```
Get-MoveRequest -Identity "TRedmond" |  
    Set-MoveRequest -BadItemLimit 20 |  
    Resume-MoveRequest
```

The first command retrieves information about the move request. This information is piped to the second command, which increases the bad item limit to 20 and pipes the data to the third command. That final command tells MRS to resume working on the move request. If things happen as we expect, MRS will complete the move request and all will be well in the world.

### **Fools Rush In Where MRS Discards Data**

Before you increase bad item limits, you must understand that MRS discards every bad item that it meets; it doesn't copy this data from the source to the target mailbox. Thus, a limit of 20 bad items means that MRS can discard as many as 20 discrete pieces of data from a user's mailbox while it's en route to a new home. This data might be unimportant, as in the case of an old calendar meeting request. Then again, it might be a message that contains an Excel attachment describing the current year's budget. Fortunately, in 99.9999 percent of cases, bad items are extremely corrupted and highly unlikely to be accessible by a client.

Letting MRS dump a lot of bad data as it moves mailboxes might be deemed a very bad thing. However, bad items are corrupted, and it's good to give the mailbox the equivalent of a colonic irrigation from

time to time, to remove accumulated debris. Indeed, Microsoft moves mailboxes frequently between databases in its Exchange Online cloud service, both to balance workload across available servers and to ensure that mailboxes stay squeaky clean and don't introduce corruption that could compromise the databases' service.

Corrupted items arise through many sources. Bugs are the most obvious and can occur on both the server and its clients. Exchange is now a very mature product, and its developers have had years to eliminate potential problems that might introduce corruption at the database level. You're likely to see more corrupted items when you move mailboxes from an Exchange 2003 server than when moving mailboxes from Exchange 2007 or Exchange 2010.

The root cause of item corruption is more likely to be a client problem than a server issue. A profusion of clients can connect to Exchange, using a multitude of connection methods, including add-on software products that integrate with clients such as Microsoft Outlook, Clashes and corruption can occur when multiple clients attempt to access the same item concurrently. Given the number of mobile devices that connect to Exchange, it should be no surprise that corruptions exist, especially in calendar items.

Based on purely anecdotal evidence in online forums, users of Research In Motion (RIM) BlackBerry devices seem to experience more corrupt calendar items than do users who access Exchange only through Outlook or Microsoft ActiveSync devices. A look through calendar items in a mailbox (use the Outlook list view) often identifies problems such as apparently blank items or items with multiple versions (i.e., conflicts). Items such as these can cause MRS problems and are best removed before mailbox moves are attempted.

The general rule is that the older the item, the more likely it is to cause a problem. If you, like me, have items in your mailbox that go back to 1996, it should be no surprise that some item properties that Exchange expects a well-behaved MAPI client to populate might not have been written correctly into the database.

## Other Bad Item Limits

Mailbox moves aren't the only work that MRS does. Any request that MRS processes can have a defined bad item limit. Think of MRS as a component that shuttles data around for Exchange, and you can understand why it implements bad item limits in everything it does. Clearly, there's no point in moving corrupted data from one place to another. It's best to detect and drop the bad stuff.

For example, if you use the [New-MailboxImportRequest cmdlet](#) to import data from a PST into an Exchange 2010 or archive mailbox, or you use the [New-MailboxExportRequest cmdlet](#) to export data from a mailbox to a PST, you'll find that you can specify bad item limits. The same zero default value is used, and you'll need to specify the `AcceptLargeDataLoss` parameter if you decide to use a bad item limit higher than 51. The [New-MailboxRestoreRequest cmdlet](#), which Microsoft introduced in Exchange 2010 SP1 as the preferred method to restore a soft-deleted or disconnected mailbox, also supports bad item limits. You might need to use this cmdlet to restore content to a mailbox from a copy in a backup-recovery database.

Importing data from a PST can be tricky. The PST file format was never designed to cater to the many ways that it's used in practice. Older ANSI-format PSTs are particularly problematic when it comes to item corruptions, and I often need to set high bad item limits (i.e., greater than 50) to successfully import data. The problem appears to be less severe with the newer UNICODE-format PSTs created since Outlook 2003 was introduced. Still, I rarely see imports of large PSTs (i.e., greater than 1GB) without encountering some corruptions.

## Rules Can Hiccup

Many users like to create mailbox rules to help automatically process incoming messages. Users like rule processing so much that they might create more than 32KB of rules in a mailbox. MRS won't process a mailbox that has more than 32KB of rules. The only solution is to specify the `IgnoreRuleLimitErrors` parameter when creating the move request or

amending an existing move request with the Set-MoveRequest cmdlet. Unfortunately, this parameter allows MRS to move the mailbox without moving any of the user's rules. The user must then recreate rules when he or she starts to use the new mailbox. Hopefully, this issue will be addressed in a future version of Exchange.

## The Good News

The good news is that a mailbox that has been successfully moved to Exchange 2010 is clean and free of corrupted items—at least, according to MRS. How long the mailbox remains clean is entirely dependent on the software that accesses and updates the mailbox. Software is getting better too, so fewer corrupted items are being created. The combination of better software and MRS sanitization should minimize bad mailbox items in the future. At least, that's the plan! ■

InstantDoc ID 142659

# Windows IT Pro

## June Schedule of Events

Web Seminars | In-Person Events | eLearning Events

### — Web Seminars —

**Wednesday, June 6, 12:00 PM**

ESG Talks About How EMC Delivers Mission Critical Confidence For SQL Server 2012

**Thursday, June 7, 11:00 AM**

BYOD and SharePoint Survey Results Webinar: What Every Organization Needs to Know

**Wednesday, June 13, 12:00 PM**

Apple Devices in the Enterprise

**Thursday, June 14, 12:00 PM**

Leveraging Subject Alternative Name (SAN) certificates to protect your applications with a single SSL certificate

**Tuesday, June 19, 12:00 PM**

Taking Touch to the Next Level: Design and Tooling Considerations for the Natural User Interface

### eLearning Events

**June 19, 21, 26, 28, July 3, 11:00 AM**

Building a Fully Functioning Windows 7 Deployment Solution

**Wednesday, June 27, 11:00 AM**

From Zero to HTML 5 in Three Easy Sessions.

# Companies that modernized with Simpana<sup>®</sup> software cut data management costs by up to **50%**

Enable faster, trimmer, better  
performing SharePoint<sup>®</sup> &  
Exchange servers

Assist upgrades with  
release independence

Simpana is integrated  
with System Center  
Operations Manager

Integrate with Azure<sup>™</sup> and  
other Cloud providers

Completely manage  
Exchange, SharePoint,<sup>®</sup>  
Windows,<sup>®</sup> Hyper-V<sup>®</sup> & Azure<sup>™</sup>



**commvault**  
**SIMPANA**<sup>®</sup>  
*software*





# Task Scheduler: All Grown Up

Leverage the power of this built-in Windows tool

**W**indows Task Scheduler is a well-known GUI utility tool for automating recurring tasks. Many people don't realize that Task Scheduler does more than schedule scripts, programs, and documents to run at specific times. You can also use the tool to display message boxes and send out email messages. What's more, you can schedule tasks to run based on a number of event triggers, including startup, logon, and even custom events. This article will highlight these event triggers and explain how to get the most from the often overlooked Task Scheduler.

## Times Have Changed

There was a time when Task Scheduler was thought to be quite basic and unconfigurable, compared with the cron jobs of UNIX and Linux systems. These days, Task Scheduler is both configurable and capable. So before looking at third-party utilities, determine whether Task Scheduler can already do what you want.

For the purposes of this article, I'll focus on the Windows Server 2008 OS. Although other OSs, such as Windows 7, might differ slightly, Task Scheduler should function in much the same way as described here.

## Task Scheduler Basics

Just as you must learn the rules before you can break them, you must become familiar with the basic operation of Task Scheduler before delving into its more esoteric uses. It won't take long.



**Rob  
Gravelle**

resides in Ottawa, Canada, and is the founder of Gravelle Consulting. Rob has built systems for intelligence-related organizations such as Canada Border Services and for commercial businesses.



**Email**



**Gravelle  
Consulting**

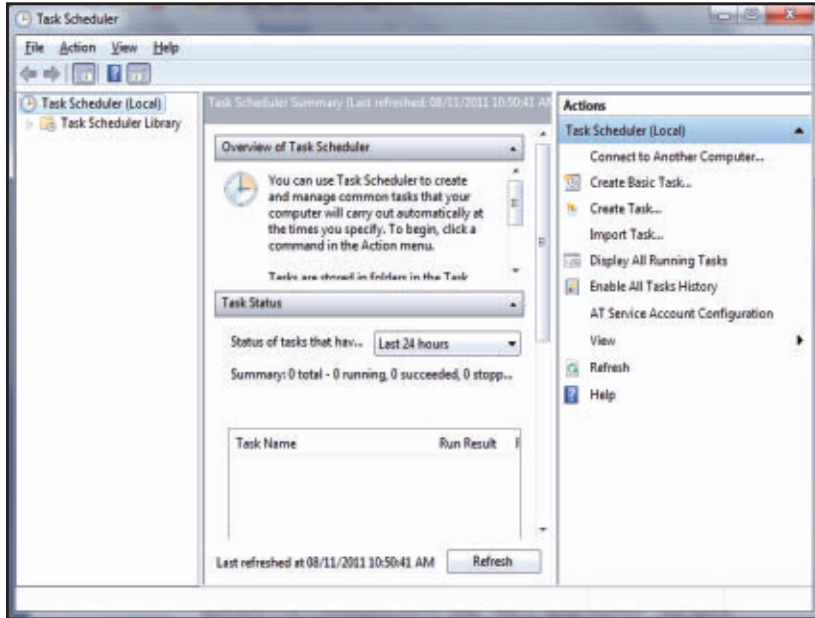


**Rob Gravelle**



To open Task Scheduler, you need administrator access. You can use the Control Panel, the Manage command, or `taskschd.msc` for this task. Figure 1 shows what Task Scheduler looks like when it first opens.

**Figure 1**  
Task Scheduler



## Creating a New Task

There are three ways to create a new task:

- Create a basic task
- Create a non-basic task
- Import a task

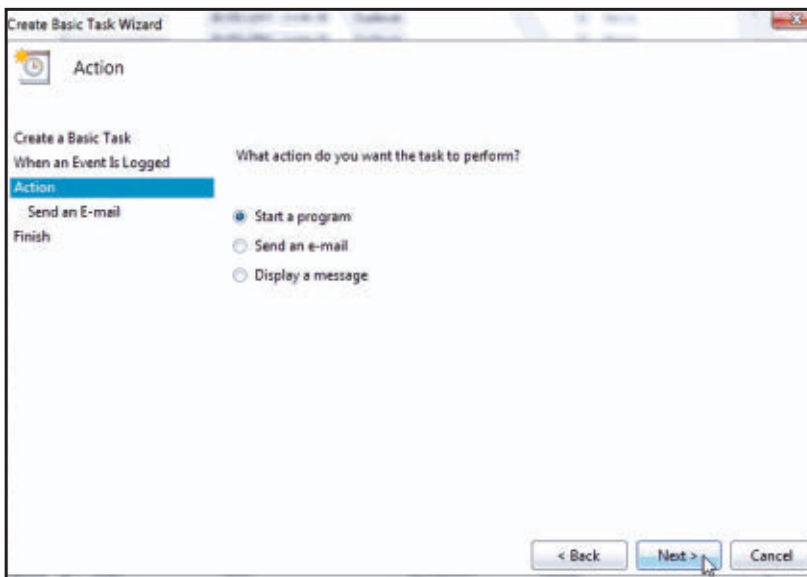
The Create Basic Task menu item will do the job for those tasks that don't require advanced settings, such as task time limits or number of repetitions. Here's how to set up a scheduled task to launch Microsoft Outlook at startup:

1. Click Action from the menu bar, and then click Create Basic Task.
2. Type a name for the task (we'll use Launch Outlook in this example) and an optional description, and then click Next.
3. You now have the option of scheduling the task to run

- on a calendar-based interval, such as Daily, Weekly, Monthly, or One Time only
- at common recurring events, such as when the computer starts or when you log on
- on specific events, such as when a specific event is logged

Select the Triggers tab and then click New.

4. From the *Begin the task* drop-down list, choose *At log on*.
5. Choose 1 minute or 2 minutes from the *Delay task for* drop-down list, to let your PC complete its startup tasks. (I'll explain more about the benefits of this decision later.)
6. Make sure that the Enabled check box is checked, and then click OK to close the New Trigger dialog box.
7. The Action screen, which Figure 2 shows, gives you three actions to choose from:
  - Start a program
  - Send an e-mail
  - Display a message



**Figure 2**  
Starting a program  
action

Select *Start a program* and then click Next.

8. On the Start a Program screen, click the Browse button to locate outlook.exe. In Windows 7, this file will be in one of the following locations:
  - For 64-bit Windows—C:\program files\microsoft office\office14\
  - For 32-bit Windows—C:\program files (x86)\microsoft office\office14\
9. Click Next to proceed to the Finish screen.
10. The Finish screen shows the task summary and a check box to display more options after the task is created. Click Finish to create the task, which then is added to the Active Tasks list.

## Running the Task on Demand

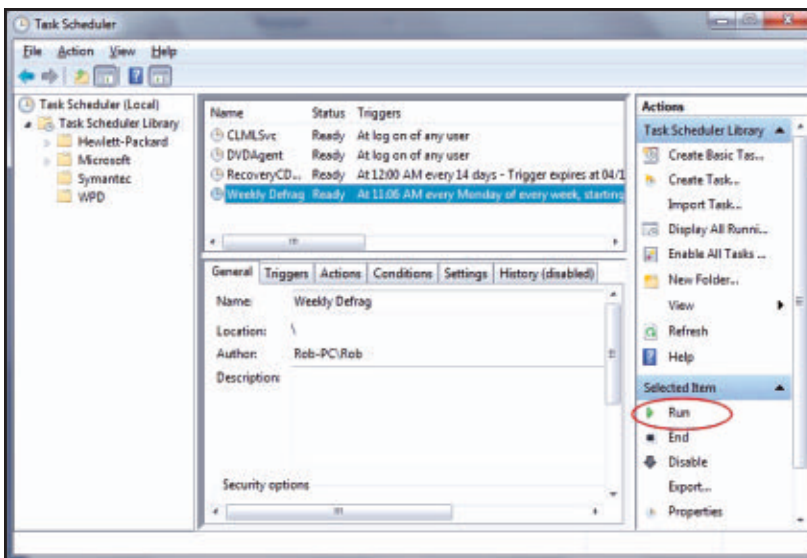
From the Active Tasks window, you can double-click any scheduled task to open its properties. Any task can be run on demand so that you can manually run it for convenience, testing, or debugging (although the latter can be built in to some extent). Here's how to run a task on demand:

1. In Task Scheduler's central pane, in the Active Tasks list, double-click an item to bring up its properties. The right-hand Actions pane shows options for the Selected Item, as Figure 3 shows.
2. Click Run to start the task. You can click the Stop button to abort the process.

## Modifying an Existing Task

You can disable a task from its Selected Item menu. To permanently remove the selected task, use the Delete button.

The center pane displays the properties for the selected task. However, these properties are read-only. To gain access to this information in an editable form, click Properties under Selected Item. This action opens a Properties window that provides more fine-grained

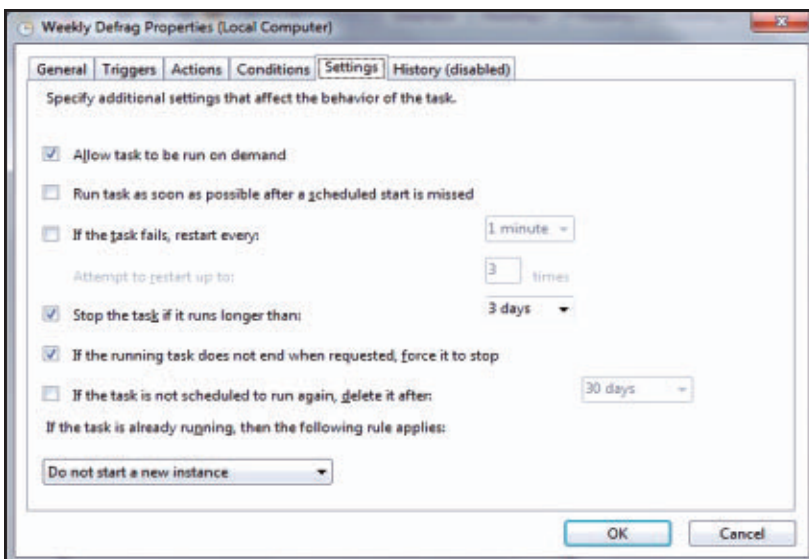


**Figure 3**  
Run command

control to the task's triggers, actions, conditions, and settings, as Figure 4 shows.

The remainder of this article focuses on these settings and how to use them to set up more intricate scheduled tasks:

- Schedule a delayed startup script



**Figure 4**  
Task Settings tab

- Schedule shutdown after the server has been idle for a certain length of time
- Execute a Windows PowerShell script
- Schedule a task based on a specific event
- Test your tasks
- Fine-tune an event-triggered task via the Task Scheduler Developer API

## How to Schedule a Delayed Startup Script

Delaying the launch of a startup script or application so that the server can finish booting is usually a good idea. There's an option for that in the Advanced Settings section of the New Trigger dialog box.

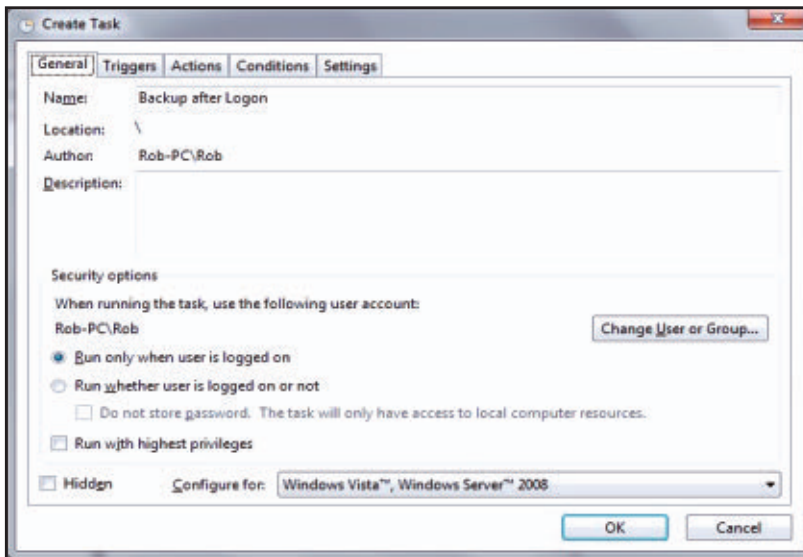
Now, you might be wondering why you would need such a scheduled task trigger when the Startup folder offers the same functionality. There are a couple good reasons. First, boot-up and logon already place a sizeable load on your computer. Starting processes directly at logon places an additional burden on your computer and can render it unresponsive for quite some time. Using a delayed scheduled task allows you to wait until the computer has finished its business. The second (and more important) reason is that scheduled tasks can be run with more rights than items in the Startup folder, which run under the logged-on user's credentials. User Account Control (UAC) requires an elevated token when attempting to perform certain tasks:

- Backing up and restoring data
- Changing group memberships
- Setting user and group security
- Creating administrative scripts
- Creating logon scripts
- Creating user and group accounts
- Deploying and upgrading software

Let's create a task to perform a backup after logon. Because we want to configure advanced settings, we'll create a regular task rather than

a basic one. The steps are much the same as the ones we used to create our Launch Outlook task, except that you'll start with the Create Task action (instead of Create Basic Task).

On the Create Task dialog box, which Figure 5 shows, you'll set the account and privileges that the task will use. With UAC, users of the admin group have two tokens. The filtered token represents standard user rights. When you click an executable and select *Run as administrator*, the full token—which contains admin rights—is used.



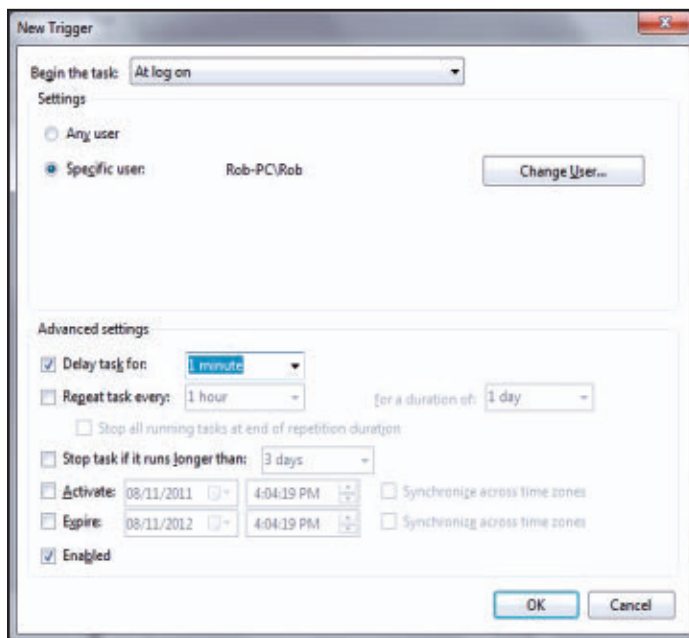
**Figure 5**  
Task Security options

Likewise, when you select *Run with highest privileges* in the *Security options* area of the Create Task dialog box, the full token is used. This setting works only when the user is in the admin group, because only those users have two tokens. To run a program with admin rights from a standard user account, you must select *Run whether user is logged on or not* and then select a user that is a member of the admin group.

Enter the username in the form *computername\username*—for instance, *server01\administrator*. Enter the password as well. On the Triggers tab, click the New button to open the New Trigger dialog box, which Figure 6 shows. The *Delay task for* drop-down list gives



**Figure 6**  
New Trigger dialog  
box



you six choices, ranging from 30 seconds to 1 day. Unfortunately, you can't define your own delay time; you must make do with the selection that is closest to what you're looking for. In our case, 1 minute should suffice. Click OK to close the dialog box.

On the Actions tab, click the New button. Click the Browse button and locate `C:\program files\cwrsync\cwrsync.cmd` in the New Action dialog box. (Note that several actions can be associated with the same task.) Click Next to proceed to the Finish screen, and then click Finish to create the task.

## How to Schedule Server Shutdown

Here's a scheduled task that uses the [Windows Sysinternals PsShutdown](#) utility to shut down a server that has been idle for a certain amount of time. This is a great way to conserve system resources while not in use.

1. Create a new task, as you did in the previous section.
2. In the Program/script text box, enter the following command:

C:\Adminutils\PSShutdown.exe

3. Enter the following in the *Add arguments* text box:

```
-s -f -c -t 10
```

4. You can choose to delay the task so that it runs only after the PC has been idle for a certain amount of time. Select 30 minutes from the *Delay task for* drop-down box.


## How to Execute a PowerShell Script as a Scheduled Task

PowerShell is the most powerful administrative tool in Windows, so it's only natural that Windows admins would want to use Task Scheduler to automate their PowerShell scripts. It turns out that these two tools are a perfect match.

You'll need to set the execution policy to execute PowerShell scripts. To execute a PowerShell command directly, specify the *command* parameter, and then use the ampersand (&), a pair of curly braces ({}), and the PowerShell command that you want to run:

```
powershell -NoLogo -executionpolicy RemoteSigned -command &{get-process>>C:\fso\mytestprocesses.txt}
```

I recommend using the Run command, as Figure 7 shows, to test your command syntax prior to scheduling your task.



**Figure 7**  
Running a PowerShell script

To execute the script, simply replace the `-command` flag with `-file` flag:

```
powershell.exe -nopprofile -executionpolicy RemoteSigned -file  
%public%\myscript.ps1
```

Note that you can include environment variables in the paths by delimiting them with percentage characters (%), as with any `Cmd.exe` command.


## How to Schedule a Task Based on an Event

Windows events are perhaps the most complicated of all the task triggers, one reason being that Windows events encompass a whole range of potential triggers. To open the Event Viewer, which is part of the Server 2008 Administrative Tools, click Start, Control Panel, System and Security, Event Viewer.

The Event Viewer, which Figure 8 shows, displays information about hardware, software, and system problems. Events are categorized by type:

- Error—a significant problem, such as loss of data or functionality (e.g., logged if a service fails to load during startup)
- Warning—not necessarily significant, but might indicate a future problem (e.g., logged when disk space is low)
- Information—information about the successful operation of an application, driver, or service (e.g., logged when a network driver loads successfully)
- Success Audit—a successful audited security access attempt (e.g., a user's successful attempt to log on to the system)
- Failure Audit—a failed audited security access attempt (e.g., a user's failed attempt to access a network drive)

You can bind an event to any of these event types by right-clicking the event in Event Viewer and selecting `Attach Task To This Event` in the



**Figure 8**  
Event Viewer

pop-up menu. This action launches the Create Basic Task Wizard, in which you can configure task options.

## Fine-Tuning an Event-Triggered Task

Since Windows Vista, Task Scheduler has contained new interfaces for C++ developers, scripting objects for VBScript developers, and a schema for defining tasks in XML. These avenues offer more flexibility than ever before for both task creation and fine-tuning.

Case in point: Suppose you want to run a task when a user is added to a specific domain local group. You'll need to hook into Audit system events, which relate to system security. Each event has a code to help identify it:

- 4730—A security-enabled global group was deleted.
- 4731—A security-enabled local group was created.
- 4732—A member was added to a security-enabled local group.
- 4733—A member was removed from a security-enabled local group.

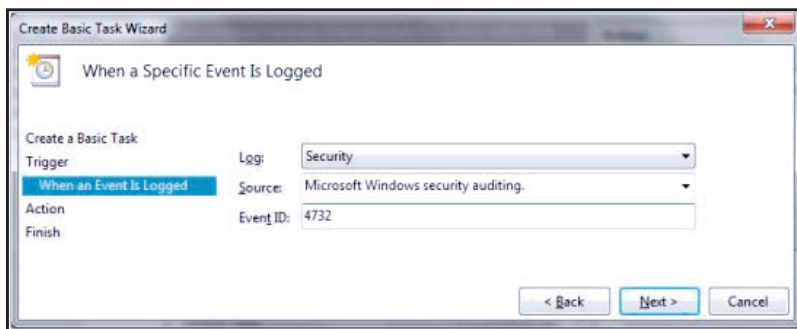
- 4734—A security-enabled local group was deleted.
- 4735—A security-enabled local group was changed.
- 4737—A security-enabled global group was changed.

Open the Create Basic Task Wizard, and then follow these steps to schedule the task:

1. Select *When a specific event is logged* on the Task Trigger screen.
2. Code 4732 identifies the addition of a member to a security-enabled local group. Enter that code in the Event ID field on the When a Specific Event Is Logged screen, as Figure 9 shows.

**Figure 9**

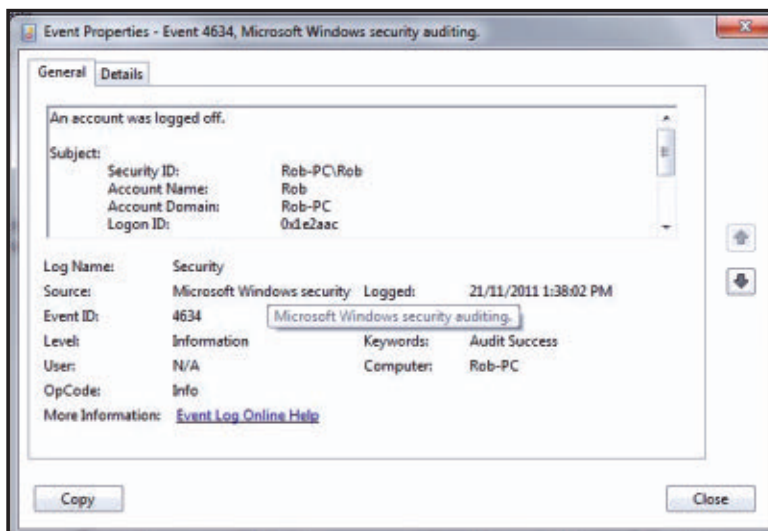
Logged event trigger details



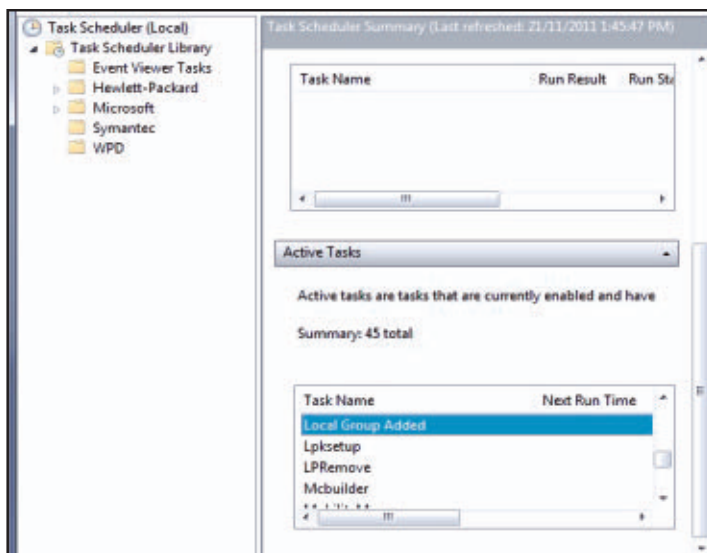
3. To find the Log and Source for an event, open Event Viewer; in the left pane, select Windows Logs, Security and open the Properties dialog box. The General tab contains both the Log Name and Source fields. The latter is liable to be truncated, but when you position the mouse pointer over the name, a tooltip appears with the full text, as Figure 10 shows.
4. On the Action screen, select *Display a message*.
5. Click Finish to create the new task.

By default, a message will be displayed every time a new local group is added. If you want to target a specific group, you need to edit the XML. To do so, you first need to export the task:

1. Double-click the new task in the Task Scheduler Active Tasks list, as Figure 11 shows.



**Figure 10**  
Determining Log and Source



**Figure 11**  
Local Group Added task

- Click Export in the Selected Item list, and save the XML file to a location of your choosing.
- Open the XML source in your favorite text editor. You'll see code that looks like Listing 1. Add the TargetUserName code to the EventTrigger section, as callout A shows.



### Listing 1: XML Code to Trigger Event on Specific Security Group Name and to Pass MemberName to Task Action

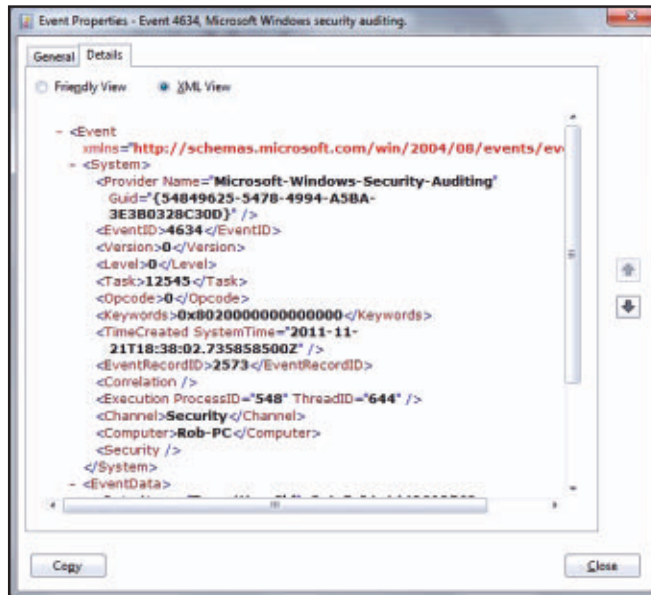
```

<Triggers>
  <EventTrigger>
    <Enabled>true</Enabled>
    <Subscription>
      <QueryList>
        <Query Id="0" Path="Security">
          <Select Path="Security">*[System
            [Provider[@Name='Microsoft-Windows-
              Security-Auditing'] and EventID=4732]]
            and *[EventData[Data[@Name="TargetUserName"]=?SECURE
              NETWORKS?]]
          </Select>
        </Query>
      </QueryList>
    </Subscription>
    <ValueQueries>
      <Value name="MemberName">Event/EventData/Data
        [@Name='MemberName']</Value>
    </ValueQueries>
  </EventTrigger>
</Triggers>

```

Now the scheduled task will respond to event 4732 only when the TargetUserName (i.e., the Local Security Group Name) is SECURE NETWORKS.

You can also filter on certain users, by using MemberName instead of TargetUserName. You can see more details about how the filter works in the XML view of an event. Simply open the event properties, go to the Details tab, and choose the XML View option, as Figure 12 shows. The full code is shown in Figure 13. You can use the data from this event in the action. To allow the EventData parameter to pass the MemberName to the action, you need to add the code at callout B in Listing 1 to the exported XML file, just before the `</EventTrigger>` line.



**Figure 12**  
Local Group Added  
task in XML view

You can now use the `$(MemberName)` variable in an action. The following code displays this variable in the message box:

```
<Actions Context="Author">
  <ShowMessage>
    <Title>Local Group Added</Title>
    <Body>The $(MemberName) local group has been added.</Body>
  </ShowMessage>
</Actions>
```

When you've modified and saved the XML, delete the original task from Task Scheduler. You can then import the modified XML. In the right pane of the Task Scheduler window, under Actions, click Import Task. Your custom task is now ready to go!

## More Than Capable

Through the manipulation of triggers, actions, and events, Server 2008 Task Scheduler offers more than ample flexibility to manage the most

**Figure 13**

Event Properties in  
XML View

```

Event Xml:
<Event xmlns="http://schemas.microsoft.com/win/2004/08/events/
event">
  <System>
    <Provider Name="Microsoft-Windows-Security-Auditing"
      Guid="{54849625-5478-4994-A5BA-3E3B0328C30D}" />
    <EventID>4634</EventID>
    <Version>0</Version>
    <Level>0</Level>
    <Task>12545</Task>
    <Opcode>0</Opcode>
    <Keywords>0x8020000000000000</Keywords>
    <TimeCreated SystemTime="2011-11-21T18:38:02.735858500Z" />
    <EventRecordID>2573</EventRecordID>
    <Correlation />
    <Execution ProcessID="548" ThreadID="644" />
    <Channel>Security</Channel>
    <Computer>Rob-PC</Computer>
    <Security />
  </System>
  <EventData>
    <Data Name="TargetUserSid">S-1-5-21-1442693562-47006223-28
      61729492-1000</Data>
    <Data Name="TargetUserName">Rob</Data>
    <Data Name="TargetDomainName">Rob-PC</Data>
    <Data Name="TargetLogonId">0x1e2aac</Data>
    <Data Name="LogonType">7</Data>
  </EventData>
</Event>

```

mundane of tasks to the most intricate. Moreover, most properties are readily available via the Task Scheduler Wizard. You can achieve more fine-grained control through the Scheduled Task API and XML schema. ■

InstantDoc ID 141796

# BitLocker in Windows 8

Changes to the drive-encryption feature enhance data protection

**W**indows BitLocker Drive Encryption is a Windows data-protection feature that Microsoft first made available in Windows Vista. BitLocker offers volume-level encryption for data stored on Windows client and server platforms. The feature protects the data when the Windows system is offline (e.g., when the OS is shut down) and can prevent data breaches such as the theft of confidential data on laptop computers.

Microsoft has continued to improve BitLocker functionality in successive Windows releases, to allow it to protect more drive and device types. In the first version of BitLocker, which shipped with Vista and Windows Server 2008, only one volume—the OS drive—could be BitLocker-protected. In Vista SP1 and Server 2008, Microsoft added support for BitLocker protection of different volumes, including local data volumes. In Windows 7 and Server 2008 R2, Microsoft added BitLocker support for removable data volumes (i.e., memory sticks and external data drives), a feature that Microsoft refers to as BitLocker To Go. In the next version of Windows, code-named Windows 8, Microsoft extends BitLocker's protection reach through support for failover cluster volumes and SAN storage.

But Windows 8 also comes with an important set of BitLocker usability enhancements that can significantly reduce the time it takes to enable BitLocker protection. These enhancements are BitLocker



**Jan  
De Clercq**

is a member of HP's Technology Consulting IT Assurance Portfolio team. He focuses on cloud security, identity and access management, architecture for Microsoft-rooted IT infrastructures, and security of Microsoft products. He's the co-author of *Microsoft Windows Security Fundamentals* (Digital Press).



**Email**

---

**Windows 8 comes with BitLocker usability enhancements that can reduce the time it takes to enable BitLocker protection.**

---

pre-provisioning, used disk space-only encryption, and standard user PIN and password selection. In this article, I provide more details about these new BitLocker features and how you and your organization can leverage them.

## **BitLocker Pre-Provisioning**

In Windows 8, administrators can enable BitLocker for a volume before the OS is installed. Microsoft refers to this as BitLocker pre-provisioning. Thanks to pre-provisioning, users can rapidly implement BitLocker protection for their data. Users don't need to wait for the encryption process to finish when they turn on BitLocker after Windows has been installed. In Vista and Windows 7, users must wait until the Windows OS has been installed, BitLocker has been enabled, *and* the entire encryption process has finished.

During pre-provisioning, Windows generates a random encryption key that BitLocker then uses to encrypt the volume. Microsoft calls the random encryption key a clear protector because it is stored on disk in an unprotected way. After Windows is installed, users can fully protect the encryption key for the pre-provisioned volume by activating BitLocker on the volume and selecting a BitLocker unlock method.

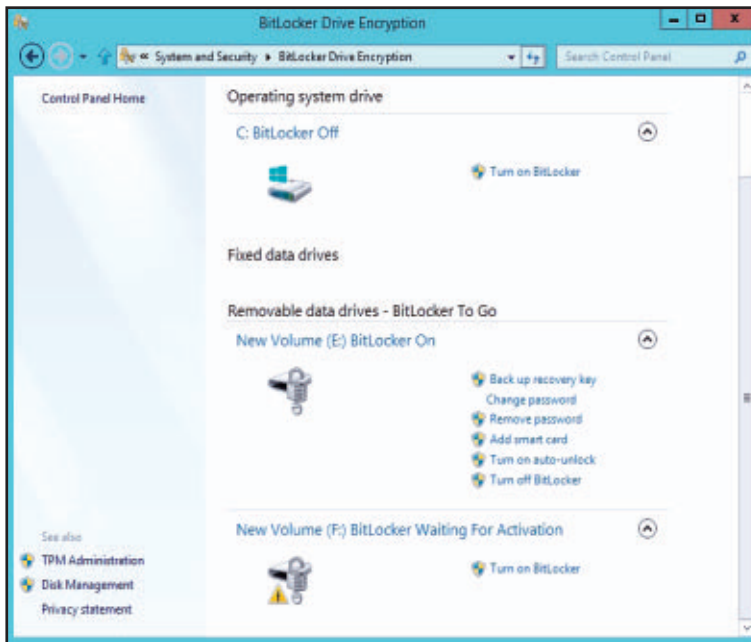
Administrators can enable BitLocker pre-provisioning from the Windows Preinstallation Environment (WinPE) by using the Manage-bde BitLocker command-line utility. WinPE is a lightweight Windows environment that is used for installing the Windows OS. For example, to pre-provision BitLocker on your F drive, type the following Manage-bde command at a WinPE command prompt:

```
manage-bde -on f:
```

Note that you need a customized WinPE image to make Manage-bde work in WinPE. (By default, WinPE doesn't include the Manage-bde tool or the Windows Management Instrumentation—WMI—objects that Manage-bde leverages.) To create this custom WinPE image, you

must add the optional WinPE-WMI and WinPE-SecureStartup components, as described in the Microsoft article “[Building a Windows PE Image with Optional Components](#).”

To support pre-provisioning, Microsoft is introducing a new BitLocker status for volumes: BitLocker Waiting for Activation. When a volume is pre-provisioned, it shows up with this status and a yellow exclamation-point icon in the BitLocker Drive Encryption Control Panel applet, as Figure 1 shows for drive F. The exclamation-point icon highlights the fact that the encryption key is still unprotected.



**Figure 1**  
Windows 8 BitLocker  
status codes

Just as you would for a regular BitLocker enablement, you can use the BitLocker Drive Encryption applet, the Manage-bde command-line tool, or Windows PowerShell BitLocker cmdlets to activate BitLocker after it has been pre-provisioned. Depending on the volume that you’re protecting, you can also choose one of the following BitLocker unlock methods:

- Five options for OS drives:
  - o Trusted Platform Module (TPM) only



- o TPM plus PIN
  - o TPM plus startup key
  - o TPM plus PIN plus startup key
  - o Startup key only
- Three options for fixed and removable data drives:
  - o Password
  - o Smart card
  - o Automatic unlock

See the Microsoft article [“How Strong Do You Want the BitLocker Protection?”](#) for a nice comparison of the unlock methods and their pros and cons.

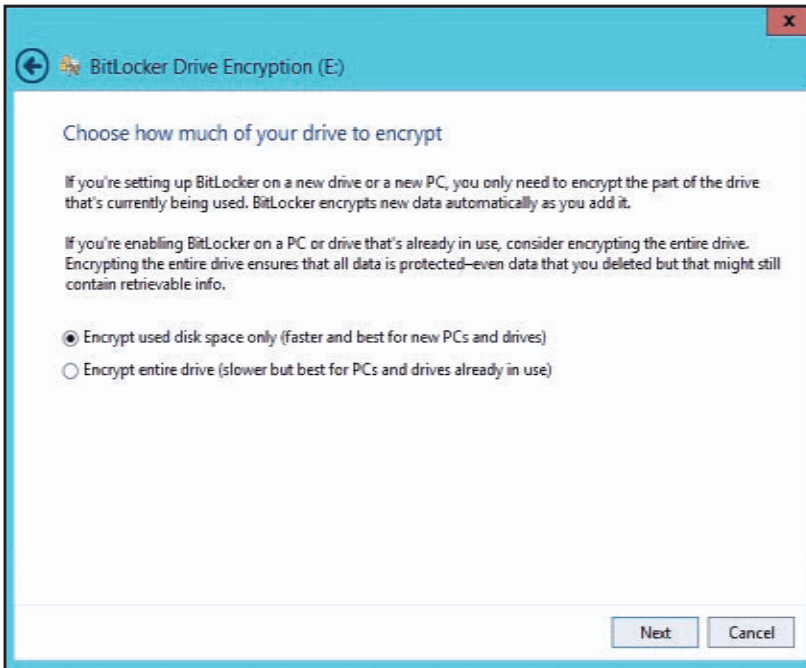
### Used Disk Space–Only Encryption

Windows 8 BitLocker supports a new encryption option that encrypts only the used space on a protected volume. Used disk space–only encryption makes the encryption of empty or partially empty volumes much faster. In previous Windows versions, BitLocker had one encryption option: encrypt everything—the data as well as all free space.

Administrators can combine used disk space–only encryption with BitLocker pre-provisioning. Enabling BitLocker on largely empty drives then becomes a process that takes seconds, and that can be invoked easily from automated Windows deployment processes and programs, by using `Manage-bde` or the BitLocker PowerShell cmdlets.

To enforce the use of either used disk space–only encryption or full encryption on domain-joined client machines, administrators can use a new set of Group Policy Object (GPO) settings in the Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption container. A new *Enforce drive encryption type* setting is available for OS, fixed, and removable data drives. These settings obviously can’t be applied to pre-provisioning; GPOs can’t be enforced before Windows is installed. If administrators don’t configure these GPO settings or set the settings to the default *Allow*

user to choose option, then users can select the encryption option in the BitLocker Setup Wizard when they turn on BitLocker protection for a volume from the Windows GUI, as Figure 2 shows.



**Figure 2**

*Choose how much of your drive to encrypt screen*


Microsoft recommends that you use used disk space-only encryption on new PCs and volumes only. Full encryption is the preferred option for volumes that are already in use. This is because free space on a used volume might still hold retrievable and valuable data, and only full encryption can ensure that everything is encrypted.

When you enable BitLocker from the command-line by using `Manage-bde` and the `-on` switch, BitLocker uses full encryption. If you want BitLocker to use used disk space-only encryption, you must add the `-usedspaceonly` switch after the `-on` switch, as Figure 3 shows.

## Standard User PIN and Password Change

A BitLocker feature that can significantly reduce and ease BitLocker deployment in Windows 8 is the ability to let a standard user (i.e., a

**Figure 3**  
Using the Manage-bde  
command-line tool



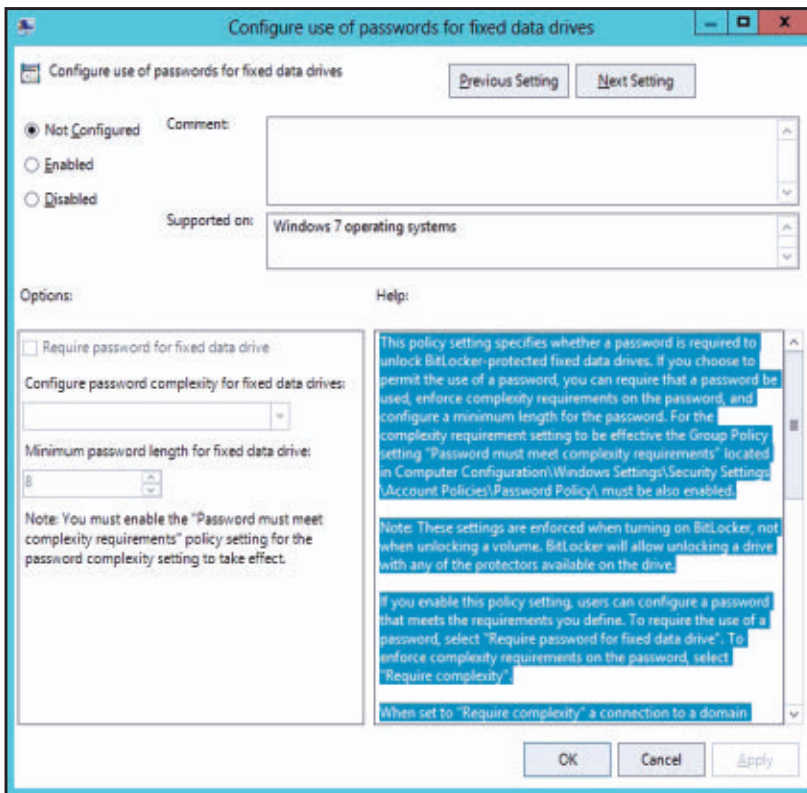
non-administrator) change the BitLocker unlock PIN (for OS drives) or password (for fixed data drives). This capability allows IT to enable BitLocker and set the same initial PIN or password on all PC images during the automated Windows deployment process. Your users can then change this initial PIN or password after the installation.

In Windows 8, standard users are entitled by default to change a volume's BitLocker PIN or password. In the BitLocker Drive Encryption applet, you'll see that the *Change PIN* and *Change password* actions aren't marked with a shield icon, as Figure 1 shows. You can change this behavior by using the *Disallow standard users from changing the PIN or password* GPO setting in the Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives GPO container. Even though this setting shows up only in the Operating System Drives GPO container, it applies to both OS and fixed data volumes.

Standard users can change the password or PIN only when they know the current PIN or password. By default, a user has five attempts to enter the correct, current PIN or password. When the retry limit is reached, the user is blocked from changing the PIN or password. The retry count can be reset to zero when an administrator resets the volume PIN or password or when the system is rebooted.

This feature also allows users to choose their PINs and passwords. Often, this capability is not to security's advantage: Users tend to use simple passwords and PINs. That's why you should always use GPO settings to enforce the BitLocker password and complexity rules. To control password complexity, you can use the *Configure*

use of passwords for GPO setting, which is available for each of the three protected drive types (i.e., OS, fixed data, removable data) in the Computer Configuration\Policies\Administrative Templates\Windows Components\BitLocker Drive Encryption GPO container, as Figure 4 shows for fixed data drives. To apply this BitLocker password-complexity requirement setting, you must also make sure that the *Password must meet complexity requirements* GPO setting in Computer Configuration\Windows Settings\Security Settings\Account Policies>Password Policy is enabled. In Windows 8, this setting is enabled by default.



**Figure 4**  
Configuring the use of passwords for fixed data drives

## Network Unlock

Network Unlock is a new unlock method for BitLocker-protected OS volumes. Network Unlock allows for the automatic (i.e., without user

intervention) unlocking of a BitLocker-protected OS volume when a Windows domain-joined desktop or server boots. In earlier Windows versions, BitLocker-protected OS volumes that are unlocked by using the combination of a TPM secret and a PIN code require an administrator to enter a PIN whenever the machine boots or returns from hibernation. This requirement makes it difficult to automatically install software and security patches on these machines.

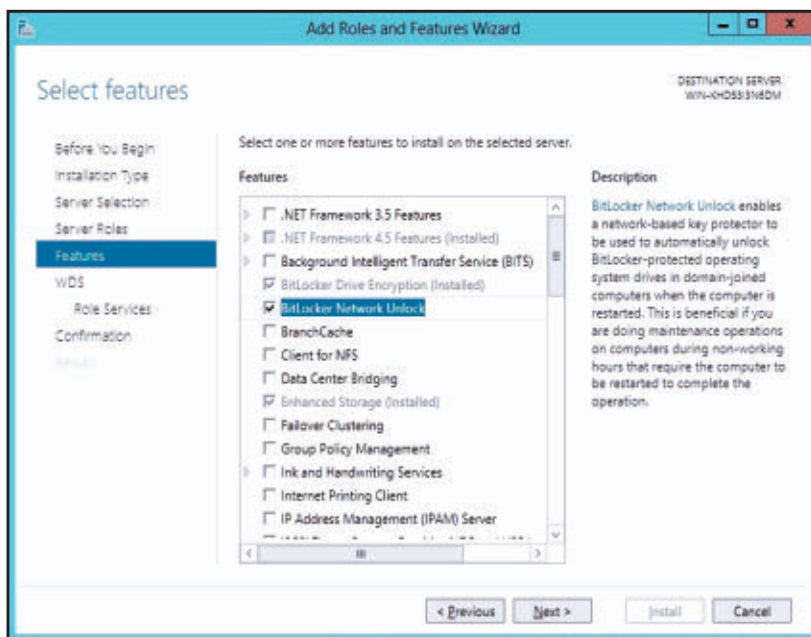
Network Unlock works like the *TPM plus startup key* unlock method. Instead of reading a startup key from a USB medium, Network Unlock uses an unlock key. This key is composed of a key that is stored on the machine's local TPM and a key that Network Unlock receives from a Windows 8 Windows Deployment Services (WDS) server on the trusted network. If the WDS server is unavailable, then BitLocker displays the standard startup key unlock screen.

Administrators can use the new *Allow Network Unlock at startup* GPO setting to control which client computers can use Network Unlock. This setting is in the Computer Configuration\Policies\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives GPO container.

The key exchange between the BitLocker-protected client and the WDS server uses DHCP. The Windows 8 WDS server role must have the optional BitLocker Network Unlock feature installed to allow the WDS server to handle and reply to the incoming Network Unlock DHCP requests. Figure 5 shows how to install the BitLocker Network Unlock option as a feature for the WDS server role in the Windows 8 Add Roles and Features Wizard.

On the client side, Network Unlock requires the client hardware to have a DHCP driver implemented in its Unified Extensible Firmware Interface. UEFI is an industry specification that defines a software interface between the OS and the platform firmware.

The WDS server also needs a special X.509 certificate and associated private key; this certificate must be present on all clients that will use Network Unlock. For more details about what the content requirements



**Figure 5**  
Adding the BitLocker  
Network Unlock  
feature for the WDS  
server role

of the Network Unlock X.509 certificate are, how to generate it, and how to push it to your clients, see the Microsoft “[Understand and Troubleshoot BitLocker in Windows Server ‘8’ Beta](#)” guide.

## Extended Storage Support

In Windows 8, Microsoft extends the use of BitLocker by enabling it to protect data on failover cluster volumes and SANs. Windows Server 8 BitLocker supports the creation of encrypted volumes on a Windows failover cluster. This applies to both physical disk resources, which can be accessed only one cluster node at a time, and cluster shared volumes (CSVs), which can be accessed by different cluster nodes simultaneously. CSV BitLocker support requires CSV version 2.0 (CSV2.0), which Microsoft introduces in Windows Server 2012.

BitLocker can now also support protected OS and data volumes that are stored on a SAN and accessed through iSCSI or Fibre Channel. BitLocker for SAN storage supports used disk space-only encryption, which is important for enabling BitLocker on large data volumes.



Finally, Windows 8 BitLocker supports a new type of disk drive that provides hardware-based encryption: Encrypted Hard Drives (EHDs). Microsoft provides an integrated interface for managing EHDs and BitLocker; this interface is basically an extension of the BitLocker Drive Encryption applet. EHDs and BitLocker each use a different approach for encryption. BitLocker protects system and data volumes by using volume-level and software-based encryption. Volume-level encryption is encryption that occurs on the volume level. Software-based encryption takes place in software.

EHDs provide Full Disk Encryption (FDE) and hardware-based encryption. FDE occurs at the disk level (i.e., on each block of a physical drive). Hardware-based encryption is offloaded to the drive's storage controller, making encryption operations more efficient.

In Windows 8, Device Manager will identify EHDs and integrate them into the OS. EHDs for Windows 8 require compliance with specific [Trusted Computing Group \(TCG\)](#) and [IEEE 1667](#) standards. You can also find more details about EHDs and about BitLocker support for SAN storage and failover clusters in "[Understand and Troubleshoot BitLocker in Windows Server '8' Beta.](#)"

## We've Come Far

BitLocker has come a long way since it was introduced in Vista. By adding new features and optimizing and refining existing features, Microsoft has significantly extended BitLocker's reach. In the meanwhile, Microsoft has gone to great efforts to make the BitLocker documentation more usable and down-to-earth. A good example is the [BitLocker FAQ](#) that Microsoft recently released. Add to this the Microsoft BitLocker Administration and Monitoring (MBAM) tool, which I discussed in my previous article "[Microsoft BitLocker Administration and Monitoring,](#)" and you can see that BitLocker is much more ready for enterprise prime time than it used to be. ■

InstantDoc ID 142661

# Integrated Network Monitoring in System Center 2012 Operations Manager

Manage non-computer devices and computer-to-device relationships

**M**icrosoft System Center 2012 Operations Manager is an upgrade to System Center Operations Manager 2007 R2. The new version of Operations Manager is evolutionary, building on a successful framework rather than reinventing too much. The upgrade offers both under-the-hood improvements and UI usability enhancements. In addition, this version of Operations Manager adds a few major features to the core Operations Manager product, one of which is a new capability for managing devices that aren't computers. A default installation of Operations Manager now includes a network-device discovery and monitoring engine that positively identifies specific network devices and graphically correlates computer-to-device relationships.

## What's New

Operations Manager can monitor hundreds of devices such as network switches, routers, and firewalls, and load balancers at the port level, and can correlate this information with server- and application-health models. Figure 1 shows the default Network Monitoring



**John Joyner**

is a Microsoft System Center MVP and senior architect at ClearPointe, a service provider that uses hosted System Center tools to deliver "24/7 EYES ON" remote management.

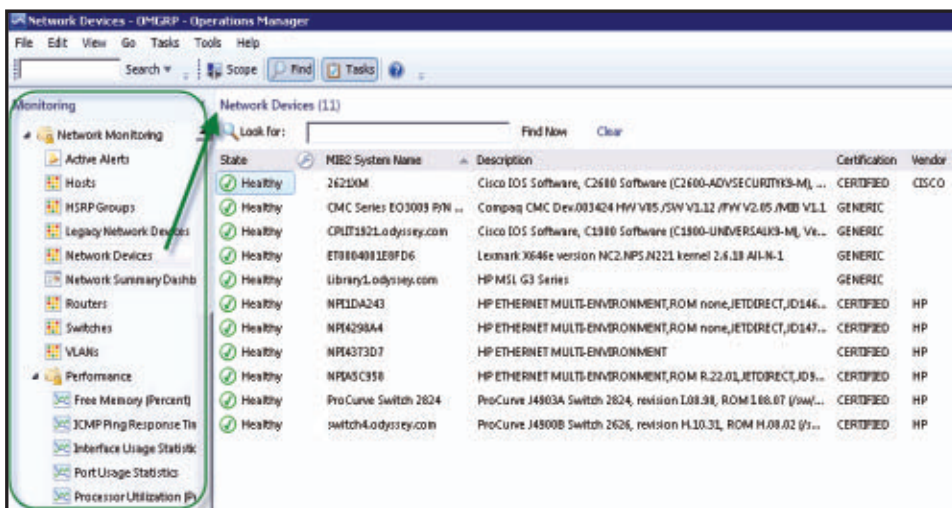


**Twitter**



**Website**

**Figure 1**  
Network Devices state  
view



views. Note that routers, Hot Standby Router Protocol (HSRP) groups, switches, and Virtual LANs (VLANs) have dedicated state views on the left. The network devices that are discovered on this small business network include a few Cisco routers and HP switches, some HP and Lexmark network printers, and a few other devices, such as an HP tape backup library. Also note the Certification column on the right; I'll cover this in more detail later.

After Operations Manager has discovered the devices on your networks, a correlation pass identifies interfaces on network devices that match up with interfaces of previously discovered devices and Windows computers. These automatically selected interfaces on network devices are monitored for performance, errors, and availability. This viewing of network-device health, in the context of computers that are interconnected by devices, is powerful and logical. The automatic selection of only key interfaces for monitoring is a clever approach that avoids collecting too much interface-performance data.

## What's Now Old

At a programming level, previous versions of Operations Manager included a rudimentary, generic network-device monitoring capability

that's based on the `Microsoft.SystemCenter.NetworkDevice` Library, which is deprecated in System Center 2012 Operations Manager. All existing third-party and in-house custom Operations Manager 2007 management packs for network devices are built on this library for SNMP device monitoring. This news isn't too bad: The earlier library was known to have scaling and performance issues, and not many IT shops have used Operations Manager extensively for network-device monitoring so far. The new version of Operations Manager includes backward-compatible support for legacy management packs written to use the older library.

The new, scalable, full-featured network monitoring in Operations Manager uses the `System.NetworkManagement` Library for SNMP monitoring. Publishers of commercial network-device management packs for Operations Manager 2007 will need to update those packs to use the new SNMP library, which includes support for the more secure SNMP version 3 (SNMPv3) protocol. Evolving beyond clear text, community string-based SNMPv1 and SNMPv2 security to encoded, cryptographic SNMP V3 security is important for confident automation of network device management.

## Deciding to Deploy Operations Manager 2012 Network Monitoring

It's impossible to manage an enterprise network by monitoring only network devices but not servers and applications. Likewise, monitoring only servers and applications, without monitoring the network devices that interconnect and support those servers and applications, is insufficient. Experienced network admins will agree that isolating intermittent or complex connectivity issues to the application or physical layer can be a time-consuming task. Any solution that integrates both layers by highlighting application-to-device dependencies is a great innovation. Such a solution speeds fault isolation and even provides input to automatic recovery workflows. Whether you decide to deploy this feature of Operations Manager might depend

on which level of integrated network monitoring you already have in place.

Monitoring both the physical and application layers of the network in a single pane of glass is a goal of most IT pros. The new Operations Manager Network Monitoring feature is Microsoft's first serious attempt to deliver that desirable, holistic picture. Many organizations today use multiple monitoring applications to instrument both the physical and application layers. It isn't uncommon for an IT shop to run Operations Manager 2007 for monitoring servers and applications, as well as running another application, such as SolarWinds Orion Network Performance Monitor or Ipswitch WhatsUp Gold Premium, for switch and router monitoring.

Where does Operations Manager Networking Monitoring fit into the management space, and should you consider deploying this feature? When making that decision, remember that monitoring network devices is not free when you use commercial software. SolarWinds charges about \$2,500 for 100 interfaces; Ipswitch charges about the same, but more generously licenses 100 devices with unlimited interfaces. The license model for network devices in Operations Manager is based on the type of network device that is being monitored. There is no charge to monitor devices that operate at network Layer 3 and lower, such as conventional switches. Devices with OS environments that function above network Layer 3 require a System Center 2012 standard management license, which costs about \$1,300. Consider the following scenarios, which might apply to organizations that use Operations Manager to monitor their networks.

***Scenario 1: Large organizations.*** Organizations with thousands of monitored network devices might already have deployed a high-investment network-monitoring solution. Operations Manager Network Monitoring isn't designed for thousands of devices like heavyweights HP OpenView and IBM Tivoli. In this scenario, consider adding Operations Manager device monitoring to speed problem isolation in specific applications. Examples of such situations include

co-monitoring of iSCSI SAN switches and network load balancers that support a critical distributed application running on Windows servers.

**Scenario 2: Midsized-to-large organizations.** Organizations with several hundred network devices might have deployed some device monitoring. In this case, take a hard look at the features in Operations Manager Network Monitoring. Can you retire an existing, secondary SNMP monitoring tool? You gain a lot with the Network Monitoring, computer-to-device correlation feature. However, you don't want to pay twice to monitor the same device. A hybrid approach might be to use Operations Manager for your datacenter core devices and to use another dedicated SNMP monitoring tool for large populations of edge switches and routers. Consider using a connector into Operations Manager, such as the SolarWinds Orion Management Pack, for those devices that aren't monitored natively by Operations Manager.

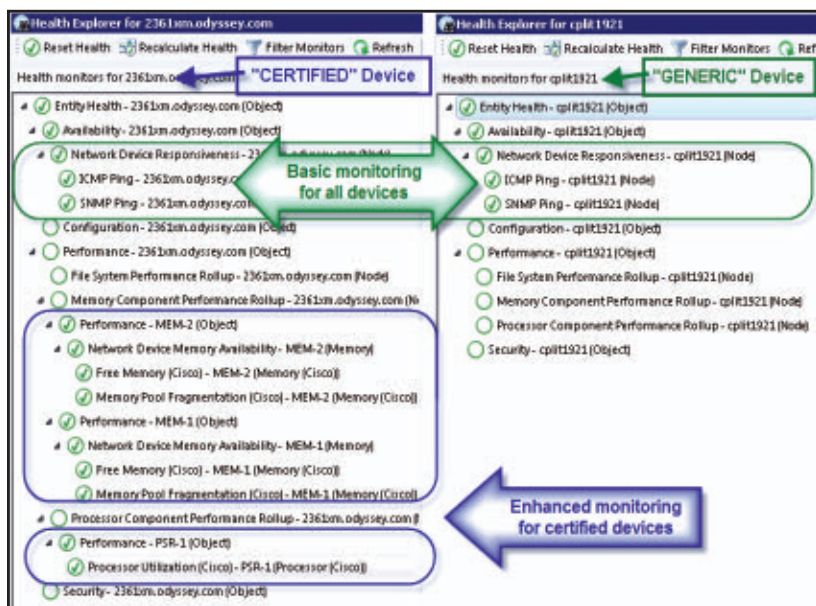
**Scenario 3: Small-to-midsized organizations.** Organizations that deploy few network-device monitoring tools might consider Operations Manager Network Monitoring for all network devices. The insight into availability metrics on your devices, and the ability to correlate network issues to server issues (without deploying any additional software), could be a big success story.

## Certified vs. Generic Network Devices

Be aware that Operations Manager Network Monitoring classifies network devices as certified or generic, depending on their status in the Operations Manager network-equipment database. Generic (or unrecognized) devices are monitored for ping or SNMP responsiveness; port monitoring looks for generic devices that support standard SNMP interfaces. Certified devices are recognized and specific additional monitoring applied. For example, the left side of Figure 2 displays the Operations Manager health model for a certified router from Cisco. This model includes monitoring of memory and processor utilization. The right side of the figure shows the health model of



**Figure 2**  
Health models of two  
network devices



a different Cisco router. This model includes only generic Operations Manager monitoring support.

System Center 2012 Operations Manager doesn't include the ability to import or compile MIB files that you supply, or to add devices to the certified database. The database of supported network devices is static and expected to be updated centrally by Microsoft. Before expecting enhanced monitoring to work with a particular model of router or firewall, test that you can monitor your key device or devices or consult the link about supported devices in the "Learning Path."

## Learning Path



### Windows IT Pro Resources

For information about rewriting legacy network device management packs:

"Migrating Operations Manager 2007 R2 Network Monitoring"

For a list of certified network devices with Operations Manager 2012 extended monitoring capabilities:

"System Center Operations Manager 2012: Network Devices with Extended Monitoring Capability"

## How to Deploy System Center 2012 Operations Manager Network Monitoring

One of the under-the-hood enhancements in System Center 2012 Operations Manager is the concept of management and gateway server resource pooling. In previous Operations Manager releases, provisioning of redundant monitoring for network devices required multiple watcher nodes against the same devices. In this version of Operations Manager, fault tolerance of monitoring nodes is automated

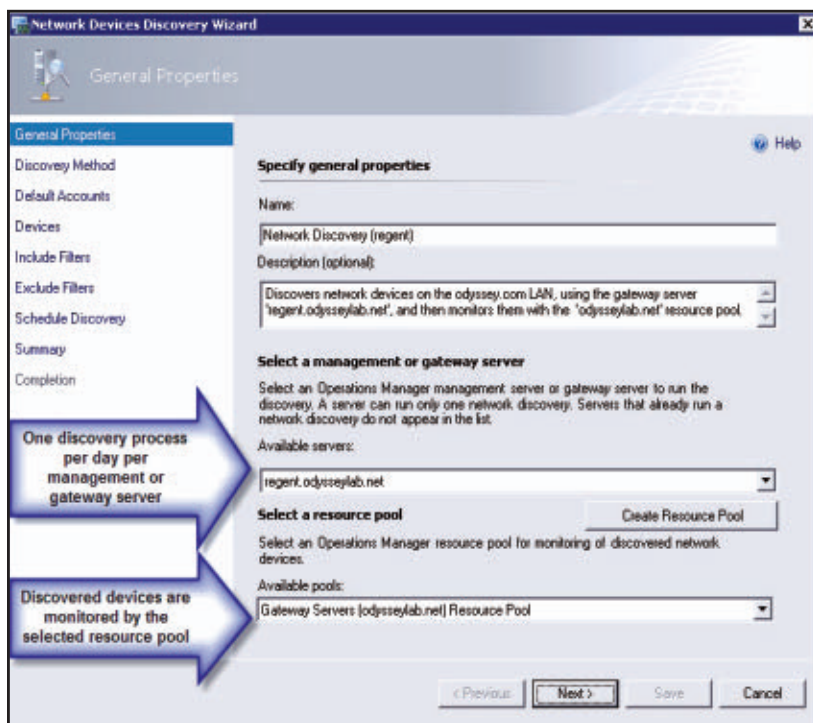
by assigning groups of managed network devices to multimember management or gateway server resource pools. In the resource pool model, two or more monitoring servers transparently load-balance and provide failover coverage for one another.

- Larger organizations (i.e., those with more than several hundred network devices to monitor) need to pay special attention to the placement and distribution of Operations Manager management and gateway servers that are members of a network-device monitoring pool. Prerelease sizing documents from Microsoft suggest that a System Center 2012 Operations Manager management group, employing two resource pools of three management servers each, can monitor about a maximum of about 2,000 network devices.
- Midsized organizations (i.e., those with up to several hundred network devices) might consider two or three servers for a dedicated (and highly available) network-device management resource pool.
- Smaller organizations (i.e., those with a few dozen network devices) can deploy the Operations Manager Network Monitoring feature on a single server, without any complications.

Your Operations Manager management group is limited to a maximum number of unique discovery rules, equal to the number of management and gateway servers in the management group. In other words, each management server or gateway server can be assigned exactly zero or one discovery rules. A discovery rule can run on the server once per day at a given time or manually only. Figure 3 illustrates how a discovery process is performed by a selected Operations Manager server and then monitored by a specified resource group.

Best practice is to consolidate all discoveries into as few rules (and servers) as possible, and to allow the automatic daily discovery process to run, optionally with the recursive discovery type selected. The intelligent process that enables both monitoring on server-connected

**Figure 3**  
Discovery and  
monitoring processes



interfaces and correct diagramming in the Network Vicinity Dashboard requires existing computers and devices to activate monitoring on discovered interfaces. Firing that discovery process daily keeps the dashboards accurate and useful, even as the server and device topology changes.

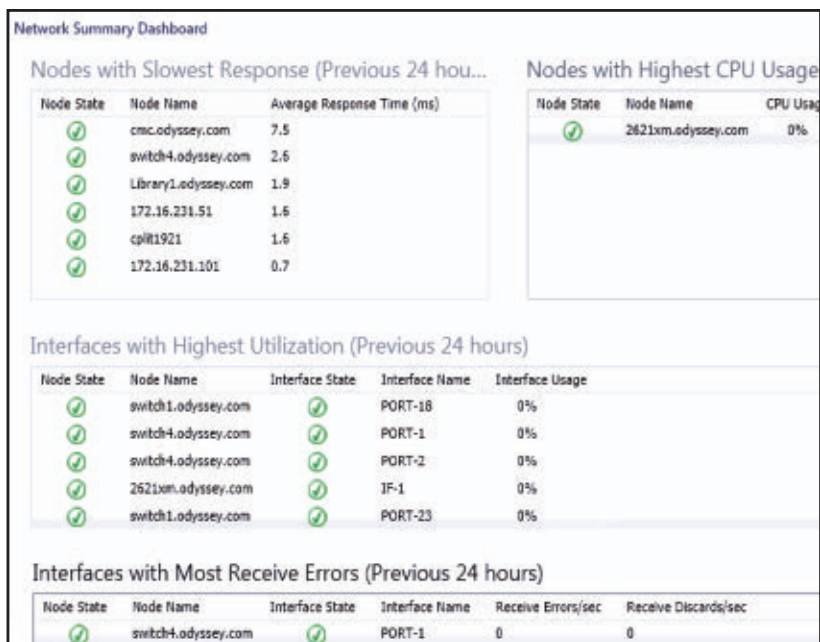
## Introducing the Network Dashboard Views

In addition to all the familiar Operations Manager view folders, such as alerts views and performance views, Operations Manager Network Monitoring introduces four new network dashboard views to convey data: the Network Summary, Network Node, Network Interface, and Network Vicinity dashboards.

**Network Summary.** The Network Summary Dashboard is the only new dashboard view exposed in the View folder hierarchy (in the navigation pane of the Operations Manager console). Therefore, it's

often the first place you'll look for a high-level overview of the health of your monitored network devices. The other network dashboards are invoked from the Network Summary Dashboard or from the task pane of any selected Windows computer or network device.

Figure 4 shows the components in the Network Summary Dashboard. These tools help you to identify the network devices and interfaces that are slowest, are busiest, or have the most errors. Use the Network Summary Dashboard to select nodes and interfaces for further analysis, then right-click the selected object or use the task pane to pivot to the Network Node Dashboard or Network Interface Dashboard.



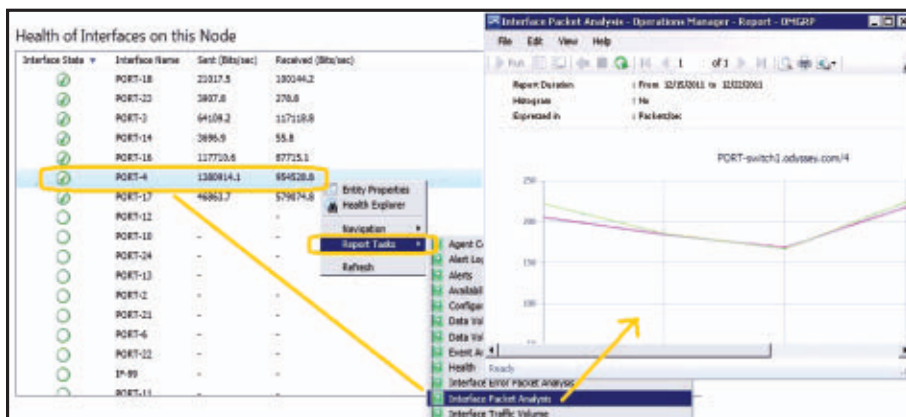
**Figure 4**  
Network Summary  
Dashboard

**Network Node.** A node is any device that connects to a network. Switches and routers are among the most common kinds of nodes. The Network Node Dashboard provides details about the health of a particular device. The upper portion of the dashboard consists of the Network Vicinity view for that node, as well as “speedometer”

gauges for node availability today, yesterday, in the past week, and in the past month. (Periods that weren't monitored are counted as available in the availability statistics so that newly discovered devices don't appear to have had outages in the gauges.)

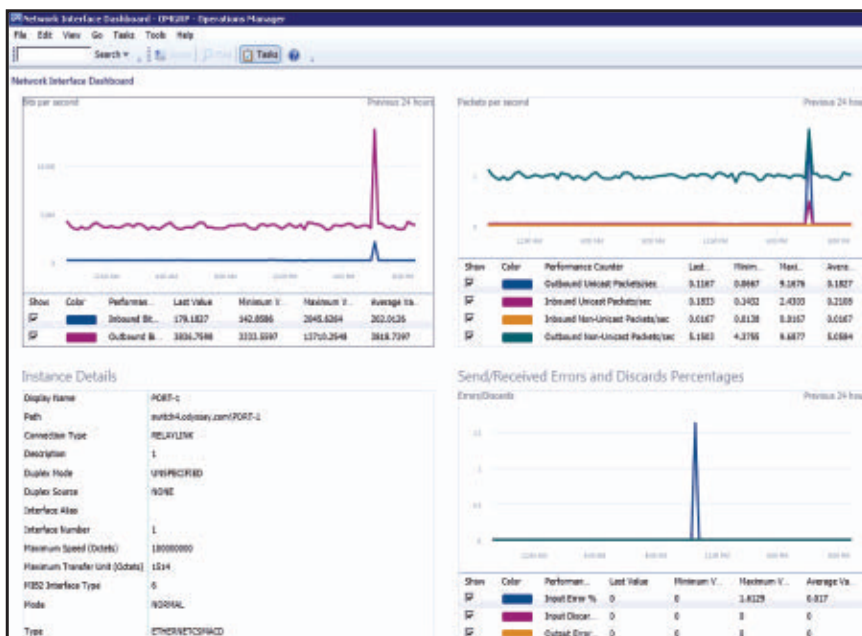
The lower portion of the dashboard includes a list of all monitored interfaces on the node. From this view, you can manually override Operations Manager's automatic selection of which interfaces to monitor. Also, by right-clicking specific interfaces, you can pivot to performance or reporting views that drill down into the near- or long-term history of an interface. In Figure 5, the Interface Packet Analysis report for port 4 on switch 1 during the previous week appears in a second window.

**Figure 5**  
Interface Packet  
Analysis report



**Network Interface.** An interface, such as a port, is a physical entity with which network connections are made. By default, Operations Manager monitors only ports that are connected to other monitored Windows computers or devices. The interface dashboard is the most detailed view of a particular interface. You can use this dashboard to zero in on a specific counter for problem investigation and capacity planning.

Figure 6 shows key counters for the previous 24 hours on a particular interface. In this case, we're looking at port 1 on switch 4, the interface that was listed in the Interface with Most Receive Errors (Previous 24 hours) section at the bottom of the Network Summary



**Figure 6**  
Network Interface  
Dashboard

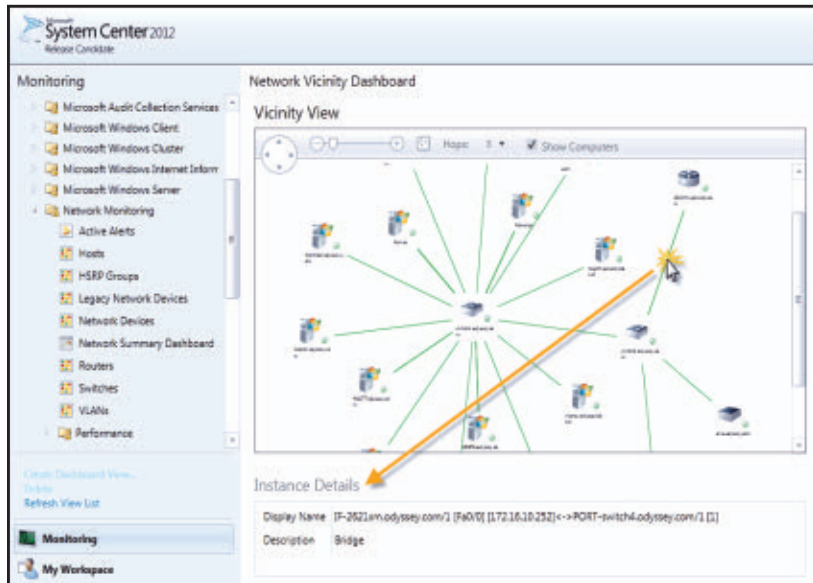
Dashboard in Figure 4. In this scenario, you get more details about the interface. Specifically, you can now answer the question, “How significant are the errors on this interface?” The Send/Receive Error and Discards Percentages chart at the lower right shows just one low spike, so the answer to that question is “Probably not very serious.”

**Network Vicinity.** Perhaps the most compelling view in the new Operations Manager Network Monitoring feature is the Network Vicinity Dashboard. This view diagrams a node, as well as all Window agent computers and other nodes that connect to that node. You can toggle up to five hops, and you can decide whether to view connected computers. Selecting a particular connection in the diagram allows you to identify which physical switch or router ports are involved; these appear in the Instance Details area of the dashboard, as Figure 7 shows.

There are some limitations in the first release of the Network Vicinity Dashboard. For one, it only works with Windows computers (not Linux computers). Second, it doesn’t take into account Microsoft Hyper-V host/guest relationships. And third, it doesn’t show network



**Figure 7**  
Network Vicinity  
Dashboard



interface teams as teamed. Another constraint is that only members of the Operations Manager Administrators group can open the dashboards, so there is no model for extending dashboard access to users who have limited-scope roles in Operations Manager.

## Closing the Gap

System Center 2012 Operations Manager adds significant new features that will be useful to many customers. Microsoft closes a gap that has existed in the Operations Manager product and makes the System Center 2012 suite more appealing. Although not a complete replacement for conventional network-monitoring tools in all environments, these features are probably sufficient, even excellent, for most small-to-midsized environments. Large organizations can instrument key datacenter devices for valuable insight into application versus physical network layer correlations that are difficult or costly to achieve with other solutions. ■

InstantDoc ID 141973



# Building a Fully Functioning Windows 7 Deployment Solution with Microsoft MVP Greg Shields

You have Windows 7 licenses, but you're not ready to deploy? Do your Windows 7 deployment skills need more than just a training class? If so, you're in the right place!

That's why in this unique five-day eLearning experience with Greg Shields, you'll get more than just training. You'll get training **PLUS** a production-ready deployment **SOLUTION** that you've built from the ground-up. In five days, you'll begin with a basic Windows Server virtual machine and finish with a fully functional Windows 7 deployment solution you've constructed yourself. Using a unique combination of labs, lecture, and online Q&A, Greg will lead you through each of the major steps in building a successful Windows 7 deployment solution. You'll learn the key steps in succeeding with your Windows 7 deployment:

- Gathering and using a hardware, software, and driver inventory.
- Building images, including every administrator's dream, *the Single Image That Installs Everywhere*.
- Incorporating deployment mechanisms using multicast, over-the-network, and USB sticks.
- Packaging applications and automatically injecting them into a Windows installation.
- Automating patches and updates to images and applications.
- Dealing with and actually fixing application compatibility issues.



Greg Shields

# SharePoint in the Cloud

How to determine whether a SharePoint cloud model makes sense



**Michael Noel**

is a partner at [Convergent Computing](#), a Microsoft SharePoint MVP, and the author of books on SharePoint, Windows Server, ISA Server, and Exchange Server. His latest book is *SharePoint 2010 Unleashed* (Sams).

Twitter



It's nearly impossible to miss all the buzz surrounding the cloud in recent years. However, there's no clear consensus on what the cloud actually is; the definition changes depending on who you talk to. This confusion exists especially when referring to Microsoft SharePoint cloud services. SharePoint can be an organization's intranet, extranet, document-management environment, records-management space, or full external website. Combined with the complexity of the SharePoint Service application layer, a cloud service could mean many different things to many different people.

In its broadest definition, the cloud for SharePoint refers to a service that allows an external provider to host and manage SharePoint functionality. This provider typically handles much of the administration and maintenance of the environment. This approach provides significant incentives for many organizations—particularly the promise of reduced overhead, a serious enticement when IT resources are consistently tasked to do more with less. At the same time, the cloud is no panacea and doesn't necessarily work for all organizations. The key to determining whether SharePoint in the cloud makes sense for your environment is first to understand how hosted SharePoint differs from an on-premises SharePoint deployment.

## What Is the Cloud?

Ask 100 SharePoint administrators to define the cloud, and you'll likely get 100 answers. Although the catchphrase is used extensively

across the industry, there is no *one* definition. There are many types of clouds, and various services offer cloud options, each of which differs slightly from the others.

For the most part, a cloud service is defined as an IT function that is housed, maintained, and (usually) administered by a third-party. The driver toward cloud services for most organizations revolves around the fact that most modern organizations prefer to focus on their core businesses rather than on the busywork of maintaining IT. This opinion is the guiding principle and main perceived benefit—mainly, the ability to offload IT to a company that specializes in it—of cloud environments.

Cloud services come in many shapes and forms: application clouds, infrastructure farms (both shared and dedicated), and private clouds. Each of these cloud options will be described in more detail in the following sections.

## Application Clouds

An application cloud is one in which the application layer of the service is offered to users. For SharePoint, this means environments that offer SharePoint sites and their associated document libraries, lists, and other application-level functionality. These types of clouds do not allow any type of low-level, infrastructure-based functionality, such as configuration of service applications in SharePoint (with some exceptions). Typically the most limited type of cloud offering, application clouds are also often the most cost-effective. The economies of scale with this type of cloud allow for massive numbers of small-to-midsized businesses (SMBs) to share the same physical hardware, reducing the total overall cost to each organization.

Several companies with SharePoint offerings can provide application clouds. These companies effectively allow you to host your SharePoint site or sites on their servers for a designated cost. This cost varies depending on the uptime requirements, level of support, amount of stored data, and provided services.

## SharePoint Cloud Vendors


This article discusses using Microsoft as your cloud vendor. There are numerous third-party options as well, including—but not limited to—the following providers:

- [Amazon Web Services \(AWS\)](#)
- [Azaleos](#)
- [CloudShare](#)
- [Connectria](#)
- [FPWeb.net](#)
- [Rackspace](#)

Microsoft provides one more well-known offering in this space, in the form of SharePoint Online, a component of Microsoft Office 365. Office 365 is Microsoft's application cloud offering and includes messaging functionality through Exchange Online, IM capabilities through Microsoft Lync, and the collaboration features of SharePoint Online. (See the sidebar "SharePoint Cloud Vendors" for a list of several third-party cloud providers.)

Figure 1 illustrates the various versions of Office 365, along with their typical list prices (subject to change). All plans include access to SharePoint Online, though some of the lower-level plans (e.g., the

**Figure 1**  
Office 365 versions





K1 and E1 plans) allow only read-only options for accessing content within SharePoint. The decision of which Office 365 plan to choose is complex and depends on the size of the organization, its messaging needs, the type of workers that it includes (i.e., knowledge worker versus kiosk worker) and whether the full Office suite of tools (e.g., Office Professional Plus) is already available or needs to be purchased. Organizations that are interested in purchasing Office 365 should check directly with Microsoft; costs can vary depending on your licensing situation and the particular geographic area in which your organization is based. Indeed, not all locations offer access to Office 365 or some of the other SharePoint cloud offerings.

## Infrastructure and Private Clouds

Infrastructure cloud options differ from application clouds in that infrastructure clouds typically allow an organization's servers to be hosted and managed by a third-party but administered and provisioned by the organization. The customer's internal network often extends into the third-party cloud provider's network so that the servers within the private cloud can communicate directly with internal servers. This approach allows the full functionality that you'd expected from internal servers but has the cloud-based advantage of outsourcing the power, cooling, and rack-space requirements.

Private clouds are often created on virtualization platforms that allow for quick provisioning and de-provisioning and a flexible systems infrastructure. For example, private cloud providers might allow the customer to provision virtual servers; the provider simply bills back the customer for the processor time, storage, and memory used. This scenario is often useful for organizations that completely virtualize their SharePoint environments. Such an environment is a supported topology that Microsoft outlines as part of its [Server Virtualization Validation Program](#) (SVVP). However, the setup does require special attention to the disk I/O, memory, and processor requirements of SharePoint Server 2010 and Microsoft SQL Server.

The main advantage that infrastructure and private clouds provide is flexibility to build servers to your specifications, to configure your farms as you prefer, and to take advantage of all the functionality of SharePoint, without needing to physically manage the servers yourself. For example, customers can take advantage of the Microsoft Business Connectivity Services (BCS) Service application to have SharePoint sites directly access data that's stored within business data repositories that are physically housed on internal database servers. (This arrangement includes scenarios such as a company that wants to display internal customer relationship management—CRM—or sales data in SharePoint sites.) This type of functionality is not readily available in the majority of application cloud models.

The disadvantage of private cloud models revolves around the need for customers to manage, maintain, patch, provision, and de-provision their servers as they would if the servers were stored internally. Much of the administration remains the same, aside from that required to maintain the virtual hosts that are stored at the virtual cloud provider.

From a migration perspective, however, infrastructure and private clouds offer the significant advantage of allowing out-of-the-box and traditional migration approaches. For example, site collections can be migrated directly via Windows PowerShell or the Stsadm tool, as opposed requiring to third-party tools.

## **Determine Cloud Requirements**

Many factors can be overlooked during the process of determining whether to move to SharePoint cloud-based solutions. Cloud environments sometimes give administrators a false sense of completeness. For example, customers assume that the cloud provider will take care of everything. This is typically not the case, however, and many organizations still must put a significant amount of thought into how the new environment will look, how to categorize content, which type of authentication to use, and other crucial decisions.

## Determine URL Strategy

A relatively minor but extremely important factor to consider for SharePoint is simply, what will the URL of the environment be? Some cloud solutions offer the ability to use a vanity URL, such as `sharepoint.companyabc.com`; others do not. For example, you might need to use a shared URL, such as `customerx.sharedproviderabc.com`. Current use of internal URLs to access SharePoint can further complicate the issue, particularly if content will be migrated to the cloud-based solution in phases. Determining what the URL will be and how to migrate content between URLs subsequently becomes a tricky task.

Some organizations deal with this issue by simply creating a landing page on the legacy environment, to direct clients to the appropriate site online. After all sites have been migrated, that landing page can then be changed to redirect content automatically. Certain application-layer intelligent systems, such as Microsoft Forefront Unified Access Gateway (UAG) or Threat Management Gateway (TMG), can also intercept and redirect specific paths within URLs.

## Define Search Requirements

SharePoint Search is a crucial but often overlooked feature within both internal and cloud SharePoint environments. The ability to easily find the content a user is looking for can make or break a SharePoint environment, so you should put proper thought into the design of Search. However, putting SharePoint in the cloud complicates this equation, as not every SharePoint cloud provider allows the ability to provide for full search across the entire environment. In addition, a search service application in the cloud can rarely crawl and index internal content, so organizations might instead decide that search capabilities need to be kept internal.

You absolutely need to know the true search requirements of SharePoint, in advance of the migration. If, for example, enhanced Search using Microsoft FAST Search Server for SharePoint is required, then you might need an internal Search instance that in turn crawls



the SharePoint cloud solution to provide federated search of both internal and external content.

### **Categorize Content**

Proper information architecture and categorization of content is key to an organization's ability to locate required content. Unless a crucial document is tagged with the appropriate metadata, it becomes invisible and there will be a significant amount of duplication of content within the environment.

SharePoint in the cloud can also pose unique challenges in that you might not have access to the Managed Metadata Service application, which allows you to apply metadata tags consistently across multiple farms. If this application isn't provided, then finding a solution that allows you to properly tag content, especially during the migration process, becomes even more crucial.

### **Determine UI Requirements**

Working on a consistent and attractive look and feel of the SharePoint environment is incredibly significant in improving end-user adoption. Users expect their collaboration environments to be easy to use, to be streamlined, and to have an appealing look.

The type of customization that is allowed with some cloud options—particularly application cloud options—might be subject to significant limitations. Some options allow customization with tools such as Microsoft SharePoint Designer, but others are completely locked down as far as what can be modified. It is important to define the requirements of the organization, in terms of UI customization, and to identify the cloud solutions that allow this level of customization.

### **Outline Security Requirements**


The security of the target environment is a crucial part of the design, particularly when dealing with cloud deployments. By their very nature, these deployments are usually accessible to the entire Internet

population. If the security of the source environment is tight and well-defined, then attempting to recreate the same security in the target environment might make sense, assuming that the same accounts or account names will be used in the cloud space. However, if the security is a mess to begin with, then restructuring security permissions during the migration might be a better option, even though it can lead to other access issues during the migration.

## Define Authentication Requirements

An important question when determining whether to use a cloud service is, quite simply, which accounts will users use to log in? Determining which accounts to use and whether users will need to enter a username and password to access the system is a major decision that has significant long-term implications.

Bear in mind that not all cloud providers offer the option to authenticate to your own internal credentials. The noticeable exception is Office 365, which allows authentication options that use your own internal Active Directory (AD) environment and Active Directory Federation Services (AD FS) 2.0. This approach, which Figure 2 illustrates, requires federation between your internal AD environment and



**Figure 2**

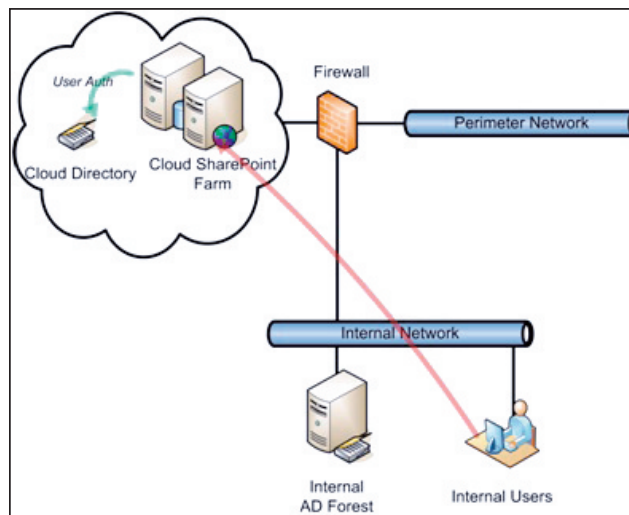
Using internal AD accounts for cloud access

the cloud provider; in Office 365, this is accomplished by using AD FS. The result is that users don't need to authenticate twice to access the SharePoint environment, their mailboxes, or the Lync client.

The alternative approach to authentication is to use the accounts that the cloud service provides, as Figure 3 shows. Cloud services typically provide this option, as does Office 365.

The obvious advantage to the AD FS option is single sign-on (SSO), but infrastructure and complexity are increased. Organizations should determine which strategy matches their requirements.

**Figure 3**  
Using accounts  
provided by the cloud  
service



## Migrating an On-Premises SharePoint Environment to a Cloud Offering

As I mentioned earlier, migrating to SharePoint in a cloud space has several limitations that can affect the migration process and limit the available options. In internal environments, an upgrade to SharePoint 2010 or a transfer of content from one farm to another supports many more options than migrating to the cloud supports. Therefore, identifying which methods are available and determining whether these methods provide the necessary capabilities is a crucial step.

## Script Migration Process

For some cloud solutions, particularly infrastructure-based solutions, you might be able to use PowerShell to script the migration of SharePoint content. PowerShell allows you to backup individual site collections from one farm and restore them to another farm. Use the following syntax to perform a backup:

```
Backup-SPSite -Identity <Site collection name> -Path <backup
file> [-Force] [-NoSiteLock] [-UseSqlSnapshot] [-Verbose]
```

You can then use the Restore-SPSite cmdlet to drop that site collection into the target SharePoint farm. The big limitation to using this approach is that it is only as granular as the site collection itself and requires a level of administrative access (on the target farm) that SharePoint cloud solutions seldom provide.

Keep in mind that PowerShell supports an export option for content, allowing a more granular export of specific document libraries or subwebs from a site collection. The syntax for export is as follows:

```
Export-SPWeb -Identity <Site URL> -Path <Path and file
name> [-ItemUrl <URL of site, list, or library>]
[-IncludeUserSecurity] [-IncludeVersions] [-NoFileCompression]
[-GradualDelete] [-Verbose]
```

The major limitation to using export commands is that they don't perform a full-fidelity extract of the contents. Indeed, export and import operations concern themselves only with content, so features such as workflows are lost during the move process. Therefore, using the Backup-SPSite cmdlet is preferable, if possible.

## Using SharePoint Designer and Outlook 2010

An alternative approach to performing migration on a small scale with SharePoint is to use client tools such as SharePoint Designer

or Microsoft Outlook 2010, which can be connected to both source and target SharePoint farms and used to manually transfer content between the environments. This scenario scales for very small environments only and doesn't work if there is a significant amount of content to transfer. The advantage to this approach, however, is that it can be used when administrative access to the back end is not provided. Another option is to use Explorer view and mapped drives to provide the ability to move content, although this approach doesn't scale very easily.

### **Third-Party Migration Tools**

Aside from these two rather limited options, the only effective way to improve on SharePoint migration is to enlist the help of third-party tools. Each tool is different, so it is important to examine each tool's functionality and compare it to the business and technical needs of the migration.

### **Understand the Pros and Cons**

Determining whether to move to a SharePoint cloud model is by no means a simple proposition. The lack of any built-in migration tools is simply one of the challenges. IT planners also need to consider the authentication options, the services that they will need, and whether security will be restructured in the process. In addition, you need to consider exactly which of the myriad types of cloud offerings to use. Understanding the pros and cons of each approach can help to streamline the decision process. ■

InstantDoc ID 142621

# Product News for IT Pros

## Brocade's ADX Load Balancer Provider Offers Cloud Optimization

At MMS 2012, Brocade touted its ADX Load Balancer Provider for Microsoft System Center Virtual Machine Manager (SCVMM) 2012, which provides real-time provisioning and elastic application delivery for cloud-optimized environments. Brocade understands the need for automation and simplified management in cloud environments. The Brocade ADX Load Balancer Provider offers seamless integration between the Brocade ServerIron ADX Series of application delivery switches and Microsoft SCVMM 2012 to enable provisioning of both application and network resources.

This integrated solution enables load balancing services to be included in the workflow for application service deployments via SCVMM, helping to streamline operations in the cloud environment and address growing application demands. The Brocade ADX Load Balancer Provider for Microsoft SCVMM 2012 is a software plug-in that enables SCVMM to natively provision and manage the Brocade ADX while making efficient use of the existing network infrastructure.

Once the Brocade ADX Load Balancer Provider plug-in is installed in SCVMM, IT administrators can incorporate load balancing services as part of an SCVMM application service template within a private cloud, assign these load-balancing services to the appropriate application service deployments, test network connectivity and Brocade ADX Load Balancer Provider interfaces before virtual machine (VM) resources go into production, and automatically configure the Brocade





ADX Load Balancer Provider plug-in to provide load balancing to specific VMs upon their deployment. For more information, check out the Brocade ADX product line at the [Brocade website](#).

### **Browsium Ion Enables IE 6 Dependent Business Applications to Run in IE 8 and IE 9**

Browsium released its next-generation solution for running legacy Internet Explorer (IE) 6 dependent web applications in modern, standards-compliant browsers. The new solution, called Browsium Ion, uses the browser engines built in to IE 8 and IE 9 to enable legacy web applications to run on Windows 7, without using the IE 6 engine. Browsium Ion builds on Browsium's first-generation product, UniBrows, and contains many new features and capabilities to put even more power, granular control, and compatibility directly in the hands of enterprise IT. Key innovations in Ion are the Adaptive IE Quirks Profile, which intelligently chooses between IE Quirks Mode and IE 7 Standards rendering, and a powerful new String Replacement feature that overrides HTML, JavaScript, or CSS from the client browser in real time with no need for server-side code changes. [Download a free Browsium Ion 60-day Evaluation Kit](#).



### **Thycotic Software's Secret Server 7.8 Keeps Enterprises' Passwords Secure**

As enterprises transfer data storage en masse to the cloud, the risk of information theft and sabotage has increased dramatically. Thycotic Software's Secret Server 7.8 password-protection software ensures the highest level of protection for organizations' most vulnerable secrets: their privileged passwords. The software suite allows controlled access to critical passwords in a central, web-based password repository. In Verizon's 2012 Data Breach Investigations Report, it was reported that over 90 percent of known enterprise data compromises in 2011 were associated with attacks that use back doors. Secret Server, by providing IT professionals the ability to securely



store, distribute, and audit password secrets, empowers organizations to secure all data from back-door breaches. For more information, visit the [Thycotic Software website](#).

## **F5 Networks' BIG-IP Helps Maximize Cloud Benefits and Productivity**

At MMS 2012, F5 Networks demonstrated its BIG-IP solutions for Microsoft private cloud deployments. The company also announced an updated F5 Monitoring Pack for System Center, which helps customers optimize resource utilization by discovering available BIG-IP devices and surfacing health statistics within System Center 2012. As administrators shift the allocation of resources, the BIG-IP system is automatically updated to ensure that the network is in sync with changes to computing and storage resources. The private cloud was the talk of the conference, and the F5 Monitoring Pack for System Center enables the key components of customers' private cloud infrastructure to work well with each other. There's a deep compatibility of System Center 2012 with F5's BIG-IP platform that can help increase efficiency and reliability while helping reduce operational expenses. Find more information at the [f5 Networks website](#).



## **Ericom Releases AccessNow 2.0**

Ericom Software announced the general availability of the new version of its HTML5-based RDP client, Ericom AccessNow 2.0. AccessNow provides browser-based access to applications and desktops running on Windows Terminal Services and RDS; virtual desktops on Microsoft Hyper-V, VMware ESX, and other hypervisors; and VDI platforms including VMware View, Quest vWorkspace, and Ericom PowerTerm WebConnect. AccessNow runs entirely within a browser and works natively with Chrome, Safari, Internet Explorer (IE), Firefox, and any other browser with HTML5 support. AccessNow is a true zero RDP client because it requires no software installation or any underlying technology on the endpoint device, such





as Java, Flash, or Silverlight. AccessNow also enables cloud hosting companies and Desktop as a Service (DaaS) providers to deliver virtual DaaS to any end-user device at any location, while supporting BYOD. For more information, visit the [Ericom website](#).

### **AppRiver Launches Secure Email-Delivery Services**

AppRiver unveiled a new secure email-delivery solution, CipherPost Pro, a service that allows businesses to easily control and encrypt email and email attachments. CipherPost Pro combines email encryption, certified email delivery, secure file transfer, and lightweight data loss prevention into one seamless offering. The product also “wraps” around any existing email infrastructure or application so organizations don’t have to replace in-place technology, including email addresses or email programs. A Software as a Service (SaaS)-based solution, CipherPost Pro enables rapid new customer deployments and is centrally managed so that IT pros can easily see the number of users, determine the amount of storage available for secure email attachments, and review usage analysis and other system data. For more information about CipherPost Pro, visit the [AppRiver website](#).



### **Advanced Systems Concepts Announces ActiveBatch Workload Automation 9.0**

Advanced Systems Concepts, Inc. (ASCI) announced version 9.0 of its ActiveBatch Workload Automation and Job Scheduling Software, which provides workload automation for the cloud computing era. The new version, expected in June, will give ActiveBatch users integrated capabilities to join both reactive and predictive forms of resource management to optimize service level agreements (SLAs) for virtually every business service throughout the modern enterprise. ActiveBatch 9.0 leverages behavioral insight such as historical workflow performance, resource availability, and capacity to govern the execution of workflows and processes by analyzing behaviors and then triggering the appropriate action. ActiveBatch then

performs predictive analysis, based on anticipated needs, to manage, provision, and schedule jobs or allocate resources in real time. For more information, see the [ASCI website](#). ■

## PAUL'S PICKS ▼

[www.winsupersite.com](http://www.winsupersite.com)



### SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

#### Nokia Lumia 900

**PROS:** Drop-dead beautiful device; large and appealing screen; excellent Windows Phone OS with visual voice mail capabilities; excellent battery life

**CONS:** Camera is surprisingly middling

**RATING:** ★★★★★☆

**RECOMMENDATION:** Microsoft's Windows Phone platform has been struggling since its inception in late 2010, but the Lumia 900, Nokia's new flagship Windows Phone handset, is a breath of fresh air and the first viable competitor we've seen to the Apple iPhone. The unibody polycarbonate design, available in black, white, or stunning cyan, features an amazing 4.3-inch AMOLED ClearBlack display that rivals anything on other platforms despite supposed limitations with its 480 x 800 resolution, and features all the goodness of Windows Phone: integration with online services, an intelligent UI with live tiles instead of static icons, over 80,000 apps, and more.

The only down side is that the 8 megapixel camera—with its widely touted Carl Zeiss optics—is lackluster, and no better than the camera on other Windows Phone handsets. Aside from the camera, the Nokia Lumia 900 is nearly perfect.

**CONTACT:** [Nokia](#)



#### Full Review

#### Windows Intune 3.0 Beta

**PROS:** Mobile device support; Active Directory (AD) integration; some Office 365 integration; new user portal

**CONS:** Expensive for small businesses; no integration with Windows Small Business Server 2011 Essentials; no software deployment for Windows Phone

**RATING:** ★★★★★☆

**RECOMMENDATION:** I've been a fan of Windows Intune, Microsoft's cloud-based PC management solution since its inception in 2010. Now Microsoft is prepping its third major version of the service, which

will manage mobile devices—Apple iPhones and iPads, Android devices, and Windows Phone handsets—in addition to PCs. That's a huge change, especially when you consider that you can use this service to deploy in-house apps to Android and iOS (but not, annoyingly, Windows Phone). Intune 3 gets deeper AD integration, a bit of Office 365 integration (though no licensing integration, nor any way to combine the services at a lower price), and a new user portal so managed users can browse for software to install rather than have it forced on them. Overall, Intune looks strong, but the \$11 per user per month fee (which now includes one PC and up to five devices each) is a bit much for small businesses. And for the beta at least, a few key features are unavailable. I expect to see the final version by Q4 2012.

**CONTACT:** [Microsoft](#)



#### Full Review

# XenDesktop 5.6



**Michael  
Otey**

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).

**Email**



**V**irtualization has become a core infrastructure component for most organizations. Server virtualization in particular has become prevalent in all sizes and types of organizations. Hosted desktop virtualization, also known as Virtual Desktop Infrastructure (VDI), is a much less widely used technology. I recently evaluated Citrix Systems XenDesktop 5.6 as a VDI provider for a Microsoft Hyper-V virtualization platform.

VDI is a misunderstood technology because it involves virtualization and desktop systems. VDI products are often confused with products such as VMware Workstation and Microsoft Virtual PC, which run the virtualization software directly on the desktop PC. XenDesktop and other VDI products aren't like this. Instead, they use a back-end virtualization server such as Hyper-V, VMware vSphere, or Citrix Systems XenServer to provide the virtualization support. The physical desktop client uses a remote display protocol such as ICA to connect to a virtual machine (VM) that's running on the back-end virtualization server. In some ways, this is much like managing a virtual server using Remote Desktop Connection. However, in a VDI implementation, the VM runs a desktop OS, such as Windows 7. Typically, there's a server between the physical desktop client and the back-end virtualization host. This server is known as the session or connection broker. Its job is to route the incoming client ICA connections to the appropriate VM image on the host. The network client then displays the desktop for that VM.

Some of the advantages of a VDI product such as XenDesktop are that it can provide centralized control of client desktops and easier migration to new desktop OSs. Centralized control reduces the number of client desktop images that you need to manage and patch as well as centralizes their location in the data center. Migration is made easier because you don't need to upgrade all your older

physical systems in order to take advantage of new client OSs such as Windows 7.

## Installing and Configuring XenDesktop

There are several components in a XenDesktop implementation. On the server side, you have the Controller, Desktop Studio, and Desktop Director components. The Controller component routes client requests to the appropriate VMs. Desktop Studio is used to create and configure collections of desktop VMs. Desktop Director is a web-based troubleshooting tool. In addition, XenDesktop requires a virtualization server. It can work with all of the popular virtualization platforms, including Hyper-V R2, VMware vSphere 4.1 and later, and XenServer 5.5 and later.

Before jumping into the installation process, you have to make sure that you have the XenDesktop requirements in place. Active Directory (AD) is required to verify the identities of the components and to allow them to communicate securely. In addition, the Controller component requires an instance of SQL Server 2008 R2 or SQL Server 2008 SP1 or later. If you don't have a SQL Server 2008 instance, the Controller installation program will install a copy of SQL Server 2008 R2 Express Edition. In order to work with Hyper-V, XenDesktop also requires either Microsoft System Center Virtual Machine Manager (VMM) 2008 R2 or VMM 2012. The VMM server must be managing the Hyper-V servers, and the VMM administrative console, with its Windows PowerShell support, must be installed on the same server as the XenDesktop server.

I installed the XenDesktop server components on a VM with 2GB of RAM and a 16GB Virtual Hard Disk (VHD) running on a Hyper-V server. For simplicity, I installed the Controller, Desktop Director, and Desktop Studio components on the same VM, but in a production environment, you would typically separate them.

The Controller component requires the Standard or Enterprise Edition of Windows Server 2008 SP2 (32- or 64-bit) or the Standard

---

**VDI is a misunderstood technology because it involves virtualization and desktop systems.**

---

or Enterprise Edition of Windows Server 2008 R2 (64-bit only). It also requires the Microsoft .NET Framework 3.5 SP1. If the .NET Framework isn't present, the setup program installs it for you.

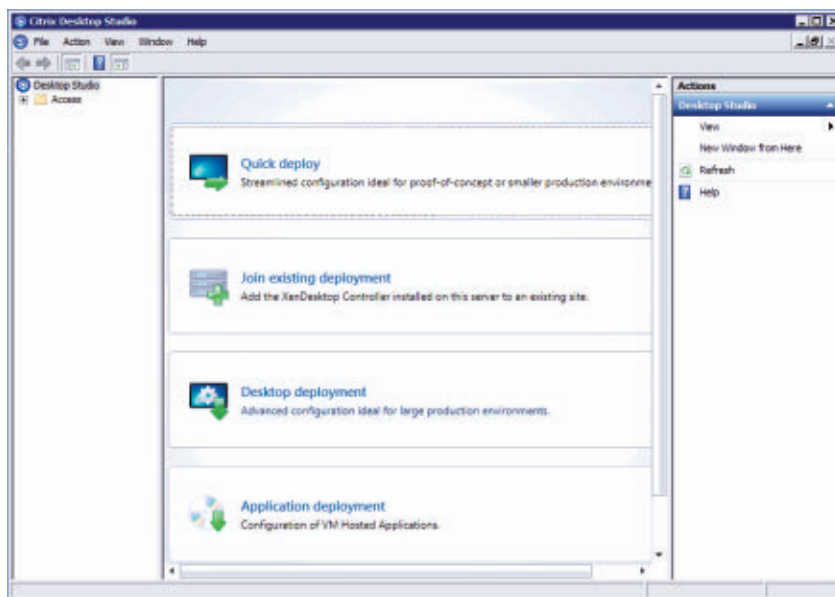
Desktop Studio supports all editions of Windows 7 (32- and 64-bit), all editions of Windows Vista (32- and 64-bit), Windows XP Professional SP2 (64-bit), and XP Pro SP3 (32-bit). It can also be installed on Windows Server 2008 R2 or Windows Server 2008 (32- and 64-bit). Desktop Studio requires the .NET Framework 3.5 SP1, ASP.NET 2.0, Microsoft Management Console (MMC) 3.0, and IIS. Desktop Director needs Adobe Flash Player.

## Preparing the Master Desktop Image and Virtual Infrastructure

After installing all the server components, I used VMM 2008 R2 to create a Windows 7 VM to act as the master desktop image. XenDesktop uses the master VM as a model for creating virtual desktops. This image contains the Windows 7 OS, the Hyper-V Integration Services components, as well as any antivirus software and other software needed by the end users. After creating the master VM, I installed the Virtual Desktop Agent on the VM and then shut it down. Next, to create the VDI infrastructure, I ran Desktop Studio.

When you first run Desktop Studio, you're presented with several different deployment options, as Figure 1 shows. The *Quick deploy* option is best for evaluation deployments. It can create up to 10 virtual desktops and perform a default configuration for the needed virtual infrastructure. The *Join existing deployment* option lets you add the Controller component to an existing XenDesktop configuration. If you're familiar with XenDesktop, you can run the *Desktop deployment* option. You use the *Application deployment* option to create virtual applications.

To set up my test environment, I choose the *Quick deploy* option. The Quick Deploy wizard asked me for the type of virtualization platform as well as the required connection information for the host and



**Figure 1**  
Desktop Studio's  
deployment options

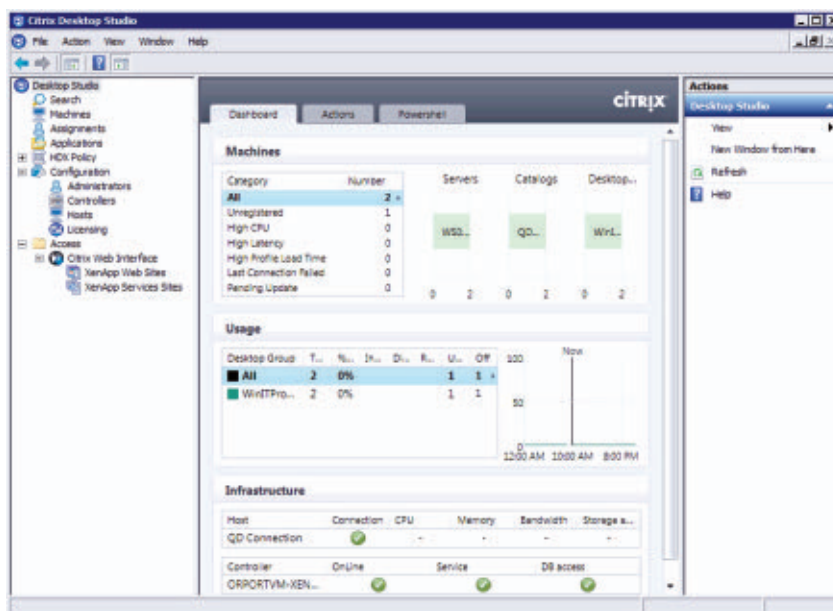
the location of the master image. For a Hyper-V implementation, this required that the Cluster Shared Volume be shared. This step had to be performed manually. Finally, I selected the number of VMs to be created and the AD users permitted to use the VMs. The wizard creates the required virtual infrastructure and creates what Citrix calls Pooled-Random desktop VMs.

XenDesktop has three basic types of desktop VMs: Pooled-Random, Pooled-Static, and Dedicated. In the Pooled-Random model, desktops are assigned randomly and, after logoff, the desktop is free for other users. Any changes made are discarded at reboot. In the Pooled Static model, desktops are permanently assigned to individual users. Again, all changes are discarded at reboot. In the Dedicated model, desktops are permanently assigned to individual users and changes persist across reboots.

After the Quick Deploy wizard completed, Desktop Studio displayed the management console that you see in Figure 2. I found the management console a bit unintuitive. Plus, the columns were too small and always needed to be expanded on my 1024 × 768 display.



**Figure 2**  
Desktop Studio's  
management console



Next, I deployed the Virtual Desktop Agent and Citrix Receiver from the installation media to the clients that I wanted to connect to XenDesktop. Citrix provides clients for the 32-bit and 64-bit editions of Windows 7, Vista SP2, and XP Pro SP3. In addition, there are Citrix Receivers for Mac OS X (Snow Leopard, Leopard, and Tiger), Apple iOS, Google Android, Research in Motion BlackBerry, and Linux. The Virtual Desktop Agents for Windows are delivered as .msi files, making it easy to deploy them using Group Policy. After installing the Virtual Desktop Client and Citrix Receiver, I needed to manually make sure ports 80, 1494, 2598, and 3389 were open. Then, the client needed to be restarted. Pointing the browser to the Controller's URL connected the client to the desktop VM that was created by the Quick Deploy wizard. Performance over my test LAN environment was comparable to a local desktop experience.

## Exploring XenDesktop's Advanced Features

One major challenge that VDI implementations face is coping with mobile users and other disconnected computing scenarios as well as

addressing the needs of knowledge workers, temporary employees, contractors, and shared workstations. XenDesktop's FlexCast technology allows it to provide virtual desktops for all of these types of users—even mobile and disconnected users. The FlexCast technology enables XenDesktop to deliver four types of desktops, which Citrix calls Hosted VDI desktops, Hosted Shared desktops, Streamed VHD desktops, and Local VM desktops.

Hosted VDI desktops are good for standard office workers. They allow personalization, and Citrix claims you can support about 150 users per server. Hosted Shared desktops don't allow personalization and are the most efficient type of virtual desktop. Citrix claims this model supports up to 500 users per server. Streamed VHD desktops use the processing power of a local client and are designed for diskless workstation implementations. This model can support 1,000 or more users per server. Finally, the Local VM desktop is intended for mobile and disconnected users. It uses the XenClient virtualization support on the target machine to run a client VM that's synchronized with a virtualization host.

XenDesktop provides three mechanisms for delivering these virtual desktop images: Installed Images, Provisioning Services, and Machine Creation Services. Installed Images are essentially Sysprep images of VMs. Machine Creation Services focuses on simplicity. This mechanism is designed to deliver a pool of dedicated VDI desktops. It's a great option for smaller organizations and can be a good way to evaluate XenDesktop. Provisioning Services provides much greater flexibility and can be used for Hosted Shared and Streamed VHD desktops.

Limited video capabilities are another concern for VDI implementations. Because the desktop VM typically runs on the host server in the data center, it's limited to the video capabilities present in the VM. Citrix addresses this problem with their High Def Experience (HDX) technology. HDX builds on Citrix's ICA protocol to provide a separate networking stream for audio and video capabilities. The

## XenDesktop 5.6

**PROS:** Thin provisioning capabilities; advanced high definition video and audio capabilities; client USB support; support for vSphere, Hyper-V, and XenServer hypervisors; support for a wide array of client platforms; support for VMM 2012; free evaluation version

**CONS:** Many prerequisites; management console difficult to use; columns in the Desktop Studio too small

**RATING:** ★★★★★

**PRICE:** \$95 per user or device, or \$195 per concurrent user for VDI Edition; \$225 per user or device for Enterprise Edition; \$350 per user or device for Platinum Edition

**RECOMMENDATION:** XenDesktop provides the best solution for implementing VDI with Hyper-V and XenServer.

**CONTACT:** Citrix Systems •  
800-424-8749 or  
954-267-3000

HDX technology enables client-side rendering of audio and video data streams enabling XenDesktop VMs to webcams and USB audio devices. Client video rendering is provided for Flash, Windows Media Video (WMV), and DirectShow. There's also server-side rendering for QuickTime and Silverlight.

### Mature, Flexible, and Capable

XenDesktop is the premiere solution for implementing virtual desktops using the Microsoft virtualization platform. It's a mature, flexible, and capable technology for delivering virtual desktops in the enterprise. The product offers a lot of options, which can make it complex. However, those options provide the flexibility needed for virtual desktop deployment at the enterprise level.

Citrix makes three editions of XenDesktop: VDI, Enterprise, and Platinum. The low-end VDI Edition includes Citrix's HDX technology but doesn't include FlexCast. The Enterprise Edition provides all the features in the VDI Edition, plus support for FlexCast and XenApp allocation virtualization. The Platinum Edition adds support for advanced features such as HDX WAN optimization, single sign-on (SSO), and SmartAccess policy controls.

I found that getting useful information from the Citrix site is a difficult undertaking. If you're interested in XenDesktop, you should try the free evaluation edition, which is limited to 10 users. You can download it from the "[Choose your XenDesktop trial](#)" web page. If you're implementing XenDesktop with Hyper-V as I did, I highly recommend that you download the Citrix white paper "[XenDesktop Hyper-V Proof of Concept Guide](#)." This guide contains essential information for this type of deployment. ■

InstantDoc ID 142617

# PowerBroker Desktops, Windows Edition

The topic of users running as local administrators seems to be on everyone's mind lately. Because antivirus and antispyware solutions often offer too little protection too late, administrators are looking for a better way to protect users' computers. BeyondTrust's PowerBroker Desktops, Windows Edition, can help you take control and put the users back in their place—as local users, not administrators.

## Installing and Configuring the Product

There are two basic components in PowerBroker: a Group Policy Management Console (GPMC) snap-in, which Figure 1 shows, and a client agent that's installed on each Windows 7, Windows Vista, or Windows XP client. Both come in 32- and 64-bit versions. The GPMC snap-in can be installed on Windows Server 2003 SP1 and later or on XP SP2 and later.

I installed the GPMC snap-in on my test domain's domain controller (DC). The only prerequisite is Microsoft .NET Framework 4.0, which must be installed separately. When you double-click the pbwd-snap32/64.msi file, a short wizard helps you install the application. It took under a minute on my DC.

I chose to install the client agent through Group Policy. By doing it this way, every time you add a new Windows 7, Vista, or XP computer to the domain and place it in the proper organizational unit (OU), the agent will be automatically installed without any other user or administrator intervention. Licensing is controlled by an XML license file that's imported into the application through GPMC.



**Eric B. Rux**

is a contributing editor for *Windows IT Pro* and the manager of technical support services at a large university in eastern Washington.

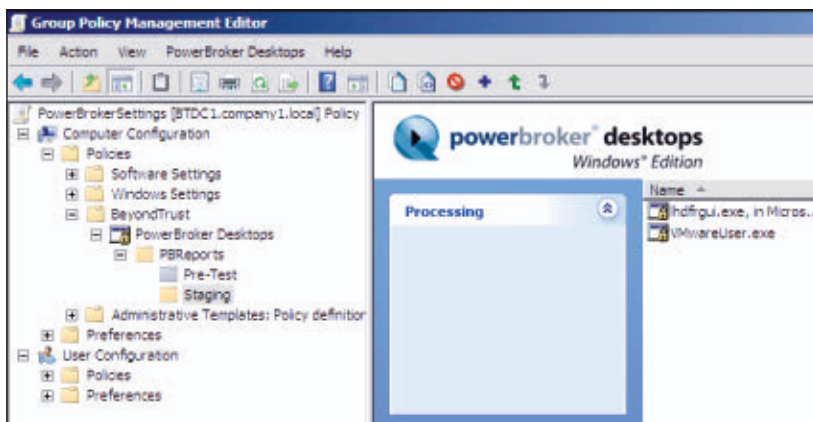


**Email**



**LinkedIn**

**Figure 1**  
PowerBroker's GPMC  
snap-in



## Installing and Configuring the Optional Reporting Solution

An optional reporting environment can be installed on a separate server. PowerBroker Desktops Auditing and Reporting requires a Microsoft SQL Server or SQL Server Express back-end database. This reporting tool uses the Microsoft Event Forwarding service to gather information about the applications that are being used and any privileges they might require. Although optional, you'll soon discover that this feature is really the heart of the application, as I explain later in this review.

Unlike the GPMC snap-in and client agent, the PowerBroker Desktops Auditing and Reporting software took much more time to install and configure. The Installation Guide does a good job of walking you through the setup tasks, but there were quite a few settings to configure.

The last step is to configure the Windows Remote Management (WinRM) and Event Forwarding on each Windows 7, Vista, or XP computer, which can be done with Group Policy. Overall, setting up PowerBroker Desktops Auditing and Reporting isn't difficult. However, the process is lengthy and prone to failure if you make a mistake or miss a configuration step.

## Using the Product

Creating a new rule that allows users to run a specific piece of software is a right-mouse click and short wizard away. In under a minute, I was

able to create a rule that allows non-administrators to use the built-in disk defragmentation tool. Instead of being met with the standard Windows 7 User Account Control (UAC) prompt, non-administrators can now easily run this tool.

The rules can be as simple or complex as you need them to be. For example, when I created the rule for the disk defragmentation tool, it didn't work the first time. This is because the rule was looking specifically for the version 6.0.6001.18000 of `lhdfmgrui.exe` from Windows Server 2008 where the rule was created. The rule didn't apply to my Windows 7 test client, as it uses version 6.1.7601.17514 of the same file. So, instead of using the file version in the rule, I used the filename and publisher (O = Microsoft Corporation, L = Redmond, S = Washington, C = US) to uniquely identify the file. After I did this, the rule worked perfectly.

If the filename and publisher aren't sufficient for security reasons, PowerBroker can uniquely identify a program by its pathname, a hash, a Windows Installer path, or an ActiveX component. Other options, such as the file location on a CD or DVD based on the serial number of the disk, are available as well.

Finally, Windows Management Instrumentation (WMI) filters are available to further define who should have the privilege of running the application defined in the rule. There are 26 filters, including filters based on whether a battery is present, CPU speed, disk space available, memory, and Active Directory (AD) security group.

When users attempt to use an application and they have the necessary privilege, the default action is to simply allow them to run it. This behavior can be changed, however, to prompt the users for justification as to why they need to use it.

## Using the Reporting Solution

Pre-authorizing applications that you anticipate that your users will need is useful, but this only takes you so far in the real world. As soon as you deploy a computer to a user, the user will undoubtedly

---

**The rules can be as simple or complex as you need them to be.**

---

## PowerBroker Desktops, Windows Edition

**PROS:** GPMC snap-in makes Group Policy integration a snap; item-level targeting lets you get extremely granular and creative when deploying rules to specific types of machines

**CONS:** Reporting structure is moderately complicated to set up; creating rules from the reporting tool is a bit clunky

**RATING:** ★★★★★☆

**PRICE:** \$30 per user; volume discounts available

**RECOMMENDATION:** If your antivirus and antispyware applications are always playing catchup, take a step back and determine why your users are becoming infected in the first place. If they're local administrators of their computers, PowerBroker can help you take back control and put the users back in their place—as local users, not administrators.

**CONTACT:** [BeyondTrust](#) • 800-234-9072

need the ability to use an application that you didn't anticipate. This is where the optional reporting functionality comes in.

The reporting application, PBReports.exe, can be accessed from GPMC or by simply creating a shortcut to the executable. If you receive a request from a user who needs to run an application, you use the reporting tool to create a rule that would allow the user to run it.

For example, suppose a user wants to run an application named MyProgram.exe. First, you use the reporting tool's query functions to narrow down the list of Windows Events to just the one that you need. You'll know that you have the right one when MyProgram.exe is listed in the Application column of the report. Next, you right-click MyProgram.exe and choose to generate a publisher, path, or hash rule. Doing so will allow you to copy XML-style data to the clipboard. This data is then copied directly into the GPMC's PowerBroker section. Although the copy-and-paste operation is all that it takes to automatically generate a new rule, the process is a bit clunky compared to other solutions in this market.

## An Easier Way to Protect Users' Computers

PowerBroker does the heavy lifting for you. Instead of having to relax NTFS or registry security for each application that would normally require local administrator privileges, PowerBroker elevates the user's privileges for just that application. The GPMC integration is super convenient, but configuring the reporting solution can be tricky if you miss one of the many steps that are required. For this reason, I'm giving PowerBroker 4 stars out of 5. ■

InstantDoc ID 142651



# SharePoint Management and Diagnostic Tool Roundup

A look at four solutions

**T**here are many tools in the marketplace that you can use to monitor Microsoft SharePoint, from free tools to home-grown solutions to comprehensive application suites. With all these choices laid out before you, it can be tough to perform a full-blown comparison of all tools. Therefore, I'll provide my opinion about the various features available in products that I've recently investigated. In this roundup, I'll discuss four products:

- Idera's SharePoint Diagnostic Manager
- ManageEngine's Applications Manager 10.3 with the Microsoft Office SharePoint Monitoring add-on
- Quest Software's Site Administrator for SharePoint
- AvePoint's DocAve 6

## SharePoint Diagnostic Manager

SharePoint Diagnostic Manager aims to deliver SharePoint administrators with a monitoring solution that's ready to go out-of-the-box, providing the information that's crucial to ensure the availability and performance of SharePoint. The metrics used are based on best practices outlined by Microsoft. Diagnostic Manager differs from Microsoft System Center Operations Manager in that only the relevant data that administrators need to keep SharePoint running smoothly is presented, making Diagnostic Manager easy to use from the get-go.



**Russell Smith**

is an independent IT consultant specializing in systems management and security, and author of *Least Privilege Security for Windows 7, Vista and XP* (Packt).



Email



Twitter

## SharePoint Diagnostic Manager

**PROS:** Comprehensive monitoring of SharePoint's back-end infrastructure; provides the relevant information that SharePoint administrators need

**CONS:** More high-level SharePoint statistics would be useful

**RATING:** ★★★★★

**PRICE:** \$1,995 per server

**RECOMMENDATION:** SharePoint Diagnostic Manager is a good choice for organizations that need more than just basic information about their SharePoint environments.

**CONTACT:** Idera • 877-464-3372 or 713-523-4433

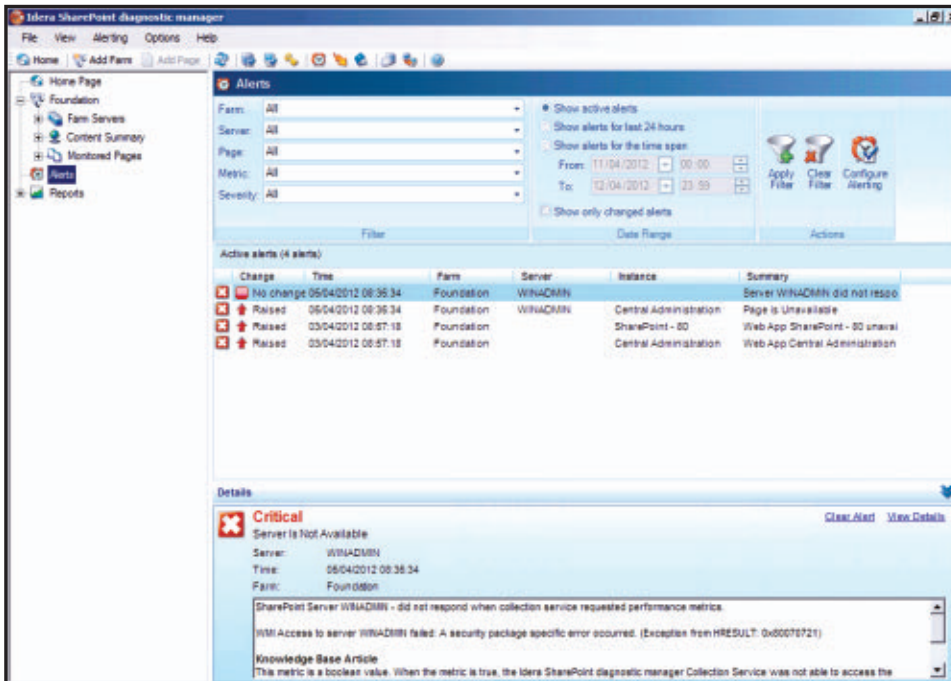
Idera specializes in monitoring SQL Server, which is the backbone of SharePoint, so it puts the company in a strong position to provide the necessary information. Diagnostic Manager pinpoints pages that load slowly (which is the most important performance indicator for end users) and provides historical data for trend analysis and forecasting.

Diagnostic Manager uses an agentless architecture with a data collection service that pulls information from the relevant SharePoint servers and stores it in a SQL Server database that's connected to the

Diagnostic Manager console. It can simultaneously monitor multiple SharePoint farms. The primary prerequisite is that the account used to run the collection service must have permission to access the Diagnostic Manager repository, the SharePoint databases, and all SharePoint servers. Diagnostic Manager supports all editions of SharePoint 2010 (including Foundation) and SharePoint 2007. The repository database must be either SQL Server 2005 or 2008. SQL Server Express editions aren't supported.

Diagnostic Manager's home page provides a nice summary of the state of the SharePoint farms, servers, and pages. It includes a common tasks area, summarizes the number of critical alerts and warnings, and provides a list of active alerts. As Figure 1 shows, when you open an alert, you're presented with all the vital information as well as additional information, such as a graphical view of the metric history or relevant information from Idera's knowledge base. Alerts can be filtered by SharePoint farm, server, page, metric, and severity for different time periods, making it easier to find information.

You can configure Diagnostic Manager to show warning and critical alerts at different thresholds for various metrics. You can also



**Figure 1**  
 Reviewing alerts in  
 SharePoint Diagnostic  
 Manager

configure an automated response to an alert, such as sending an email, writing to the event log, or generating SNMP traps.

Diagnostic Manager monitors SharePoint servers and processes, providing data on servers' states, active alerts, resource usage, search performance, indexer status, Microsoft Excel calculation performance, and more. When viewing SharePoint information, a Microsoft Office-style ribbon lets you change the view from a dashboard to more detailed information. Some of the views include Performance Monitor-style counters.

Page monitoring is especially important to understand the performance as perceived by end users. Page load time data is collected periodically. Diagnostic Manager can monitor page performance against different web front-end servers if required. Page component analysis shows rendering speed for HTML controls, web parts, and web controls.

In addition to monitoring SharePoint servers, Diagnostic Manager monitors IIS, showing important information such as Active Server

Pages (ASP) requests per second and current versus maximum connections. Diagnostic Manager also monitors SQL Server, displaying database size, log file size, and fragmentation.

Diagnostic Manager contains a large collection of reports out-of-the-box to help you diagnosis problems, present capacity planning data to management, and design projects. You can also create your own custom reports.

Diagnostic Manager is an impressive tool that Idera is continuing to develop to provide administrators with more SharePoint-specific information. The UI is well organized, is intuitive, and responds quickly. Despite the huge amount of information available in Diagnostic Manager, it's really easy to focus quickly on what's important and establish the root cause of any problems.

## Applications Manager with the Microsoft Office SharePoint Monitoring Add-On

Applications Manager uses SharePoint APIs and doesn't connect directly to back-end servers. As a result, the monitoring capabilities of Applications Manager are somewhat limited compared to a fully featured solution such as SharePoint Diagnostic Manager.

Applications Manager has four dashboards, one of which is the Default Dashboard, which Figure 2 shows. It gives an overview of all the services being monitored, including Application Managers' own components (e.g., MySQL, Tomcat). The Infrastructure Snapshot section shows system availability and health, which are graphically displayed with a green or red dot. If there's a red dot, you can drill down to a root-cause analysis window, which shows what the problem is and the root cause.

### Applications Manager 10.3 with the Microsoft Office SharePoint Monitoring Add-On

**PROS:** Easy to set up and configure; monitors more than just SharePoint

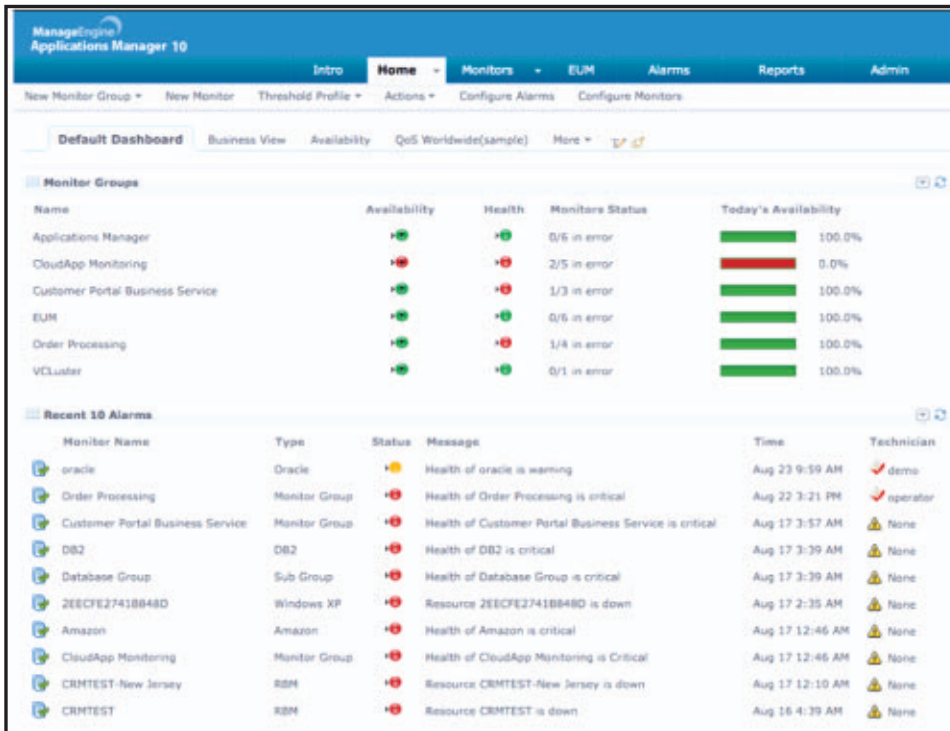
**CONS:** Limited SharePoint monitoring

**RATING:** ★★☆☆☆

**PRICE:** Pricing starts at \$795 for ManageEngine Applications Manager plus a \$245 flat fee for the SharePoint add-on

**RECOMMENDATION:** This solution is best used for situations in which a large number of different applications need to be monitored, but not in any depth.

**CONTACT:** [ManageEngine](#) • 888-720-9500

**Figure 2**

Reviewing the services being monitored in Application Manager's Default Dashboard

Applications Manager can monitor a wide range of different applications and systems, many of which are open source. There are built-in thresholds and anomaly profiles for SharePoint monitoring. Plus, you can create your own. Alarms can be set for when a threshold is exceeded or an anomaly is detected. The Alarms tab lets you drill down for more information about each alert, including a history. You can configure an alert to trigger an action, such as sending an email, sending an SMS message, or restarting a Windows service. You can also configure service level agreements (SLAs), which are displayed graphically so you can easily identify underperforming monitors.

The Microsoft Office SharePoint Monitoring add-on for Applications Manager monitors a relatively small set of metrics for Office search, Excel services, SharePoint services, document conversions, ASP, memory, and the web cache. Although this monitoring will be enough to gather basic information on SharePoint availability and

performance, it's probably not sufficient for diagnosing and finding the root causes of complicated problems.

Applications Manager comes in two editions: Standard and Enterprise. The Standard Edition is limited to monitoring 250 applications, whereas the Enterprise Edition comes with additional scalability and failover support. It's very easy to get up and running, but Java must be pre-installed.

## Site Administrator for SharePoint

Site Administrator for SharePoint's information portal is a little less streamlined than those in the Idera and ManageEngine solutions. The Site Administrator portal focuses on information growth and the number of SharePoint sites, documents, lists, and so on. Site Administrator lets administrators browse SharePoint and gives relevant information on objects, which is useful for capacity planning.

As with Applications Manager, Site Administrator doesn't expose information about the back-end servers that keep SharePoint running. So, while it's possible to create custom reports from Site Administrator's information repository, the data is limited and specific for certain tasks.

Quest purchased a product named Security Explorer, which is now integrated into Site Administrator, for discovering and managing the security of SharePoint objects. The search functionality is very powerful and can be used to discover which users have access to SharePoint documents—something that would be hard to establish without a specialized tool. It's also possible to perform security maintenance tasks, such as granting, revoking, duplicating, and reassigning user permissions, as Figure 3 shows.

The metrics include information on subsites, lists, and document versions, which are represented

### Site Administrator for SharePoint

**PROS:** Ideal for monitoring information growth and managing security

**CONS:** Limited administrative capabilities; expensive

**RATING:** ★★☆☆☆

**PRICE:** \$2,995

**RECOMMENDATION:** Site Administrator for SharePoint is suitable for organizations that want to plan for future storage growth and need a SharePoint security management solution.

**CONTACT:** Quest Software • 800-306-9329 or 949-754-8000



**Figure 3**

Assigning user permissions in Site Administrator for SharePoint

numerically and in graphs that show information growth over a configurable time frame. On the metrics screen, objects are ordered showing the biggest on disk first but can be expanded to show more if necessary.

Site Administrator can be set up to monitor multiple SharePoint farms. The HTML interface works fairly quickly, but occasionally the web back-end displayed errors instead of the required content. This was usually rectified by a simple browser page refresh.

## DocAve 6

DocAve integrates a series of SharePoint management applications (i.e., modules) into one portal. In this portal, administrators can perform a range of tasks, including performing backup and restore operations, managing content, optimizing storage, and running infrastructure and usage reports.

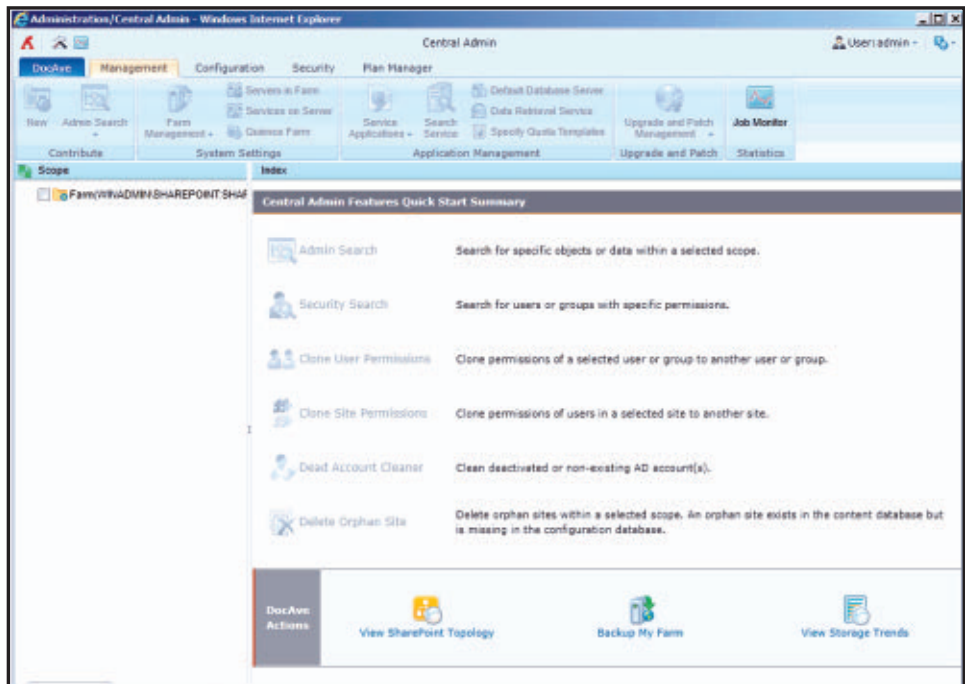
Installing the DocAve server component was easy. There's the option to use a built-in database or a SQL Server database. One possible downside to DocAve is that it requires an agent to be installed



on SharePoint to collect data. Nevertheless, installing the agent is also easy. When the server component is installed, you're supplied with a passphrase that's needed when you install any agents that will connect to the server. Both agent and server installs make proper prerequisite checks to ensure that the installation will complete successfully.

The server administration software itself is Silverlight-based. If you can live with that, the administrative GUI is cleanly presented in an Office-style ribbon interface, as Figure 4 shows. Despite the tidy interface, I found the GUI less intuitive to use than the GUIs in SharePoint Diagnostic Manager and Applications Manager.

**Figure 4**  
Working in DocAve's  
GUI



DocAve offers some neat modules, such as Connector, Backup and Restore, and Content Manager. Connector allows users to work with documents hosted on network file shares as if they've been imported into SharePoint. Connector reduces costs by enabling organizations to store data outside of SQL Server databases, which can be expensive to manage.

With the Backup and Restore module, you can back up your SharePoint environment (including all web applications, databases, front-end web elements, and content) and restore it at a farm, database, or granular level. I started by trying the module's Granular Backup feature. I was directed to a window listing a series of necessary steps with links. A simple wizard would have been welcome to avoid the need to constantly refer back to this list. The first step is to create a storage policy. The provided link opens a new window, which frustratingly can't be resized, which might be a restriction of the Silverlight back end. At the top of the window, I had to click again to actually create the storage policy. During the process, DocAve prompted me to create a logical device to use for storing the backup, which was handy. Notwithstanding my gripes, almost all backup software is non-intuitive and requires something of a learning curve.

The Content Manager module provides a Windows Explorer-type interface divided vertically into two panes so you can copy or move content between SharePoint servers. I found that this worked well, although browsing the content tree was a little sluggish. With a large SharePoint site, this might possibly prove to be an exasperating experience. Although the UI animations are slightly irritating, I liked the inclusion of Test buttons, which let you test what you're configuring.

I really liked the storage optimization features in DocAve, especially Connector and Archiver (which moves expired data out of SharePoint to a cheaper storage location). Support for enterprise storage solutions, such as the Dell DX Object Storage Platform and Data Archive's EMC Centera, is included out-of-the-box. Integration with Microsoft Office 365 is a bonus, and DocAve can be used to manage hybrid onsite/cloud SharePoint infrastructures.

## DocAve 6

**PROS:** Comprehensive SharePoint management tool; modular solution

**CONS:** Clunky Silverlight interface

**RATING:** ★★★★★☆

**PRICE:** \$3,995 per web front end (SharePoint Administrator); other components' prices depend on configuration

**RECOMMENDATION:** DocAve is an extensible platform for all SharePoint needs. If you have a large SharePoint installation to manage, I recommend looking at some of the clever modules in this package.

**CONTACT:** AvePoint • 800-661-6588 or 201-793-1111

There's a comprehensive set of reports built into the product, although without the option to create your own. The DocAve suite packs in so many features that it's impossible to cover them all here, but AvePoint seems to have all bases covered. If you have a large SharePoint installation to manage, I recommend looking at some of the clever modules in this suite.

## SharePoint Toolset

Before making a decision about which product is best suited to meet your organization's needs, you should look at each one in a lab environment and carry out your own testing. I found SharePoint Diagnostic Manager to be a comprehensive monitoring solution for SharePoint that should meet the needs of even the most complex SharePoint environments. The GUI is pleasant to work with and provides out-of-the-box views of all the information that will be important to SharePoint administrators without the need for understanding what elements of the back-end infrastructure need monitoring.

The SharePoint add-on for Applications Manager is a light-weight solution that can't provide information on SharePoint's SQL Server back-end infrastructure, which makes it unsuitable for larger SharePoint installations. However, Applications Manager is easy to use and monitors a wide array of services that make up the crucial IT infrastructure in an organization.

Site Administrator is a tool that's designed for two specific functions: storage capacity planning and security administration. The pairing of these two functions seems a little odd but might be exactly what some organizations are looking for.

DocAve is a comprehensive and extensible platform for managing SharePoint. Despite my reservations about the Silverlight administration console, it's hard not to be impressed with the range of solutions built into DocAve. ■

InstantDoc ID 142826

# Computer-Based Training for Enterprise IT Systems

Tips on how to evaluate CBT products

**M**ost likely, you're managing a diverse set of software. To keep up with new software, you need new approaches. Like most IT pros, you need training, but taking time away from the office for live training is tough and expensive. Your boss might be demanding certifications to keep your job, or maybe you're seeking a performance bonus. Perhaps you just need additional training for a new project. Whatever the reason, getting certified and continually learning new technologies is hard work but can lead to improved performance and greater professional glory. This buyer's guide provides some tips on how to evaluate computer-based training (CBT) products for popular enterprise IT applications from vendors such as Microsoft, Citrix Systems, VMware, and Cisco Systems.

As shown in the buyer's guide table, many training products go beyond the boring text-based tutorials you remember from high school. Video training has revolutionized the training industry. It's now a major force in preparing for certification, obtaining crucial job skills, and keeping up with the frequent releases of new products. The formats are diverse: online learning, downloadable content, intranet, and even application appliances preloaded with multiple courses.

I'll show you how to make the best decision on how to spend your precious money and training time. After all, if you plan to spend 20, 30, or even 50 hours watching videos, you'll want to ensure you're getting the maximum value from your investment.



## Tony Bieda

has over 10 years of experience as a systems engineer. He's an MCSE and MCDBA. His areas of focus are network infrastructure and Windows system architecture.



Email



LinkedIn

---

**If you plan to spend 20, 30, or even 50 hours watching videos, you'll want to ensure you're getting the maximum value from your investment.**

---

## Important Considerations

When choosing a CBT product, you should first scan the primary objectives and review the lesson plans to determine whether they meet your certification or general learning goals. It's helpful to note the subtopics and time devoted to each segment to ensure it focuses on your weaknesses or specific topic requirements. If there are prerequisites, review those courses as well. It can really add to the time and financial commitment if you need to take classes in addition to the one you targeted.

To ensure mastery of the subject, most videos are taught by instructors with product-specific certifications, but some are taught by experts with the highest levels of certifications. One subjective item is the instructor's content-delivery and teaching style. The objectives might be rock solid, but if you feel the instructor is boring or doesn't adequately explain the concepts, you're less likely to watch the videos or receive value from them. You can review videos many times, so it's important to find some enjoyment in watching them. To check the instructor's style, watch a few sample videos of the instructor (most sites allow this). If no samples exist, ask the company if you can receive a sample.

Some vendors offer additional course delivery options (e.g., online, CD-ROM, downloadable file, intranet, appliance). If you're pursuing multiple certifications or have multiple learning goals, consider an annual or time-based subscription. If multiple users are obtaining training, using online or intranet training might be less expensive and offer you additional options for reviewing more course content. Purchasing an appliance or a full online training library could be thought of as a companywide IT learning portal, but check the licensing agreement carefully to ensure how many people can concurrently access the content and if named users are required. Be sure to monitor subscription end dates so time doesn't run out before everyone has completed the training. Extensions can be purchased, but that would increase the costs.

If you're frequently watching the content in areas where internet access is poor or nonexistent, such as on trains or planes, consider

electronic downloads. DVDs and CD-ROMs are convenient methods used by most vendors.

In the past, many vendors had very simple interfaces, in which starting, pausing, and stopping videos were the major features. Now vendors are adding additional features, including the ability to print. For example, AppDev bundles built-in printable books and TrainSignal includes printable instructor notes. Indexing works especially well for exam reviews, on-the-job skill reviews, and fast access to concepts. The ability to create notes is helpful to save key thoughts or information. Bookmarks are crucial for noting areas for quick reference in the future, such as difficult concepts or key findings.

Hands-on labs can be approached in different ways. A popular approach is to base them on a fictitious company, which lends itself well to working through examples on a test system during the instructor's video. However, if examples build on each other, it can be more difficult to set up sample scenarios if you choose to scan the material.

When making your selection, note how often course updates are provided. Frequent course updates are necessary to ensure the content is accurate and available for the latest versions of the product. Check the fine print to understand if product upgrades and new versions are included with the purchase price, or if there is a low or no cost option for obtaining new content. Costs vary, so it's crucial to prioritize the features you would like.

The cost per hour of training isn't a good judge, and neither is cost per package. Cost should be considered in the context of the quality and suitability of the training material. To save money in a corporate environment, be sure to ask about multiuser discounts.

## An Excellent Way to Build Your IT Skills

For IT pros, students, or those looking to break into the IT industry, CBT is an excellent way to build your IT skills at a pace you prefer. You're only limited by your time and desire to learn. ■

InstantDoc ID 142763

Company	AppDev	CBT Nuggets	K Alliance	LearnKey
Series name (if different from company name)				
Licensing costs	\$395 to \$795; call for multiuser pricing	\$24 to \$1,999	\$499 to \$1,999	\$380 to \$580 for most single courses; series prices vary
Does price include updates?	Contact vendor	Yes, as long as customer's subscription is valid	Free updates for 12 months	Contact vendor
Are courses taught by certified instructors?	Yes	Yes	Yes	Yes
Video formats available	Windows Media Video (WMV)	DVD, electronic download, MP3, streaming	FLV, MP4	WMV
Are videos recorded based on final versions of software?	No	Yes	Yes	No
Is course content regularly updated?	New product releases are listed as released	Yes; timelines are product-specific	Yes; 4 to 8 months based on demand	Many products are current versions (check website for more details)
Course types	CD-ROM, DVD, online learning	CD-ROM, downloadable file, online learning	CD-ROM, intranet, online learning	CD-ROM, intranet, online learning
Are the courses geared toward specific certifications or job skills?	Certification, plus on-the-job skills	Both	Job skills	Primarily certification, but also job skills
Are pre- or post-assessments included?	Yes	Yes	No	No
Do the course objectives follow the certification objectives?	Yes	Yes	Yes	Yes
Is the content based on case studies?	No	Yes	No	No
What courses are available for professional certifications?	Microsoft SharePoint, Microsoft SQL Server, many Microsoft developer courses, Windows Server 2003 and earlier	Microsoft, VMware, Citrix, Cisco, CompTIA, IT Infrastructure Library (ITIL), Information Systems Audit and Control Association (ISACA), Check Point, EC-Council, Help Desk Institute, MySQL, Oracle, Project Management Institute, Juniper, Linux	Microsoft, Cisco, Oracle, and Red Hat	Microsoft Exchange, SQL Server, Windows Server 2008, Cisco
Included features	Built-in book you can print, subtopics, table of contents, bookmarks, searching	Indexed for searching, table of contents, note taking, bookmarking, MP3, forum community, supplemental material	Searching, bookmarking, print, note taking	Glossary, indexing, links to supplemental information and other resources
Are any hands-on labs included?	Yes	Yes	Yes	Yes



Company	Microsoft	Microsoft	Microsoft	Microsoft
Series name (if different from company name)	Microsoft Press	e-Learning	Jump Start	Microsoft Certified Master (MCM) SQL 2008 Exam Prep Videos
Licensing costs	Free to \$69.99; subscriptions through Safari Books Online	Approximately \$20 per hour of training	Free	Free from TechNet
Does price include updates?	eBook updates are free when purchased through O'Reilly Media	Free updates for 12 months	N/A	No cost for updates
Are courses taught by certified instructors?	Yes	No	Yes	Yes
Video formats available	N/A	Flash, Silverlight, WMV	AVI, iTunes, WMV, Zune	Contact vendor
Are videos recorded based on final versions of software?	No	Yes	No	Yes
Is course content regularly updated?	N/A	Updated as needed	Typically retired, as this is early readiness	Yes, with each Microsoft product release
Course types	Appliance, CD-ROM, downloadable file, eBook library, online learning, printed book	Downloadable file, intranet, online learning	Downloadable video, intranet, online learning	Online learning
Are the courses geared toward specific certifications or job skills?	Both	Yes; IT pro, developer, and information worker	No	Yes
Are pre- or post-assessments included?	Yes	No	No	No
Do the course objectives follow the certification objectives?	Yes	Yes	No	Yes
Is the content based on case studies?	Yes	No	No	Yes
What courses are available for professional certifications?	Microsoft	Microsoft Certified IT Professional (MCITP), Microsoft Certified Solution Developer, Microsoft Office Specialist	Microsoft Windows Server 2008 R2, Server Virtualization	Microsoft Certified Master (MCM): SQL Server 2008
Included features	Print books, eBooks, practice tests, lab instructions	Full text indexed for search, table of contents, progress tracking, limited print functionality, glossary	Videos and PowerPoint presentations only	N/A
Are any hands-on labs included?	No	Yes	No	No

Company	SkillSoft	TestOut	TrainSignal
Series name (if different from company name)		LabSim	
Licensing costs	Contact vendor	\$495	\$97 to \$397 per course for a single-user license
Does price include updates?	Contact vendor	Free updates for 12 months	No
Are courses taught by certified instructors?	No	Yes	Yes
Video formats available	Web based	Silverlight Smooth Streaming	AVI, iPod, VideoMP3 Audio, WMV
Are videos recorded based on final versions of software?	No	Yes	Yes
Is course content regularly updated?	New releases appear on website (update cycle not available)	Updated as needed	Yes; we make it a priority to update our training content as quickly as possible
Course types	Online learning	Online learning	DVD, online learning
Are the courses geared toward specific certifications or job skills?	More certification, but also job skills	Certification, with hands-on virtual labs that increase job skills	Both; courses build hands-on skills that are immediately applicable on the job and they provide comprehensive coverage of exam objectives
Are pre- or post-assessments included?	Yes	Yes	Yes
Do the course objectives follow the certification objectives?	Yes	Yes	Yes
Is the content based on case studies?	No	No	No
What courses are available for professional certifications?	Cisco, Microsoft	MCITP, Cisco Certified Network Associate (CCNA), A+, Network+, Security+, Linux+, Systems Security Certified Practitioner (SSCP)	Microsoft, Computing Technology Industry Association (CompTIA), VMware, Citrix, Cisco, and more
Included features	Localized content for various languages, scenarios	Searching, bookmarks, note-taking ability, print content, closed captioning, print video transcripts, dictionary, list of objectives, reporting	Table of contents, PDF of instructor notes, instant online access, multiple file formats
Are any hands-on labs included?	Yes	Yes	No

# No User Wants to Migrate Like a Wildebeest

## How to Avoid the Perils of The Great Migration to Exchange 2010

Looking to upgrade your email environment to Exchange 2010 or migrate it to the cloud? Concerned about the size of the task or the risks involved?

Join Karl Sand of Binary Tree and Julian Martin of Mimecast for an interesting take on the challenges of migration. Check out this session to learn how to eliminate risk and complexity prior to migrating, and how you can streamline and simplify your migration.

[Click here](#) to register for the webinar.

**Watch this video to see a clip from the session!**



**BinaryTree** | **WindowsITPro**

# Insights from the Industry

## Privacy Is the New Security

Internet Explorer (IE), through tracking protection lists, can protect you from websites that have a fairly wanton approach to your personal information. However, tracking protection lists aren't enabled by default and require a bit of technical know-how to configure. They certainly aren't something that your average computer user knows anything about. That's something that leaves me flummoxed, especially given the recent dispute about Google's treatment of Platform for Privacy Preferences (P3P) on IE. Why is this IE privacy feature so hidden, when privacy is so important?

IE always notifies me about unsafe sites and asks whether I want to disable add-ons. Why doesn't it notify me by default when a site such as Facebook or Google wants to use cookies across multiple websites? "By default" is important here—because if the behavior isn't by default, only knowledgeable nerds are going to be able to find and enable it. Protection or privacy isn't something that should depend on technical knowledge; it's something that should happen automatically. This setting should be something that you can override if you choose to expose your secrets. Your activity shouldn't be automatically tracked by third parties just so they can show you more relevant advertisements.

It's pretty clear that voluntary standards aren't going to cut it when it comes to privacy on the Internet. Enforcing rigorous privacy standards goes up against some lucrative business models. Changing

default settings on IE 11 or IE 10 to make them privacy-first browsers would provide users with protection in the same way that IE 9, IE 8, and IE 7 increased protection against malware. Explain to your friends that their social networking site can and does track them to all those sites with free videos of naked couples' calisthenics, and their response is unlikely to be, "Well, I'm okay with that."

Actual people consider privacy a part of their security. The people who argue otherwise generally belong to organizations that make a lot of money by tracking people across the Internet. Don't look at the man behind the curtain! The major alternative browsers to IE (Firefox and Chrome) are both funded primarily by Google (Chrome directly, and more than 90 percent of Firefox's funds come from Google). There's no incentive for either browser to be the first to make the protection of a web surfer's privacy the core trait of its browser. Do you think Facebook would make privacy a priority if the company funded a browser? Firefox and Chrome offer privacy as an add-on, but they won't run by default because doing so would kill the goose that lays the golden eggs.

As to whether or not users want this functionality: What do you think the result would be if a dialog box popped up asking users if they wanted to let Facebook or Google track their activity across multiple websites? I'm guessing that users would click the "Are you kidding?" button if it were available.

—Orin Thomas  
InstantDoc ID 142309



## Orin Thomas

is a contributing editor for *Windows IT Pro* and a Windows Security MVP. He has authored or coauthored more than a dozen books for Microsoft Press.



**Email**



**Blog**

## Proactive Data Management for E-Discovery

Have you ever had to respond to an e-discovery request for your company? Sometimes just the thought of facing legal action can make seasoned Microsoft Exchange Server administrators want to turn in their passwords and retire or find a new career as a janitor. Exchange

Server is a trove of information in almost any organization, but will you be able to find what you need when slapped with a discovery request?

Although email and Exchange Server aren't necessarily the only targets of e-discovery, it's quite possibly where the bulk of your data resides. "Email makes up so much of what is discoverable," said Barry Murphy, co-founder and principal analyst for [eDJ Group](#), an analyst firm focused around e-discovery and information management. "There's this notion that everything passes through the Exchange server at some point, so it's usually the number-one priority in any kind of e-discovery," Murphy said.

No one knows better what data lives in Exchange than the Exchange administrators who deal with it daily. Therefore, it makes sense that those admins should take an active—or indeed, proactive—role in protecting that data and ensuring the company is ready to face any legal challenge. That doesn't mean you need to know what's inside every email message on the system—which of course would get you into a whole different kind of trouble—but at least you should know where the data is, and try to maintain only the data that's required.

The introduction of Personal Archives in Exchange Server 2010 as well as the continued push to move messaging data to the cloud can only complicate the matter of managing your data. And at the same time, data protection legislation isn't going away—if anything, it's getting stronger. "It's an incredibly complex world of data," said Kevin Foisey, chief software architect at [STEALTHbits](#). "As you impose more and more rules over the top for the management of that data, the frameworks that do the management of that data need to be more and more agile."

So what can you do to manage data? A few things spring to mind. First, if your organization is in a highly regulated sector, such as health care or financial, make sure you're up-to-date with the latest data security procedures and rules that apply. Even if you're not part of such an industry, familiarize yourself with the data protection legislation

for your state or region. Wherever possible, take it as a best practice to adhere to these standards even if they might not apply to your specific organization: Make your data safer than what's required.

Second, if you haven't already done so, consider implementing a comprehensive data retention policy for email and other documents within your organization. This process is where you set rules for how long an email message, for example, will be kept in a user's Outlook before being automatically deleted; and it's also where you define what type of documents must be maintained for longer periods. With Exchange Server 2010, you can [apply retention tags and policies](#) to automate this process; other Exchange versions and mail systems will have their own versions of this procedure.

The important thing here is to work with your company's legal department up front to determine what policies are appropriate and how long different types of data should be kept. Users tend to want to keep everything in their Inbox, but by getting rid of the unnecessary and useless bits, you'll not only be using less storage but also have less data to search if you do get slapped with an e-discovery request. You could even argue (and yes, I would, as part of a company that's gone through this process—see “[Establishing an Email Retention Policy: The Legal Perspective](#)”) that forcing users to hold on to less can ultimately make them more productive. As part of partnering with legal, make sure expectations are clearly communicated to end users because you don't want them just squirreling away email messages in local PSTs or other places to prevent them from their proper expiration.

As a third preventive measure, you can look at what third-party products can do to help you secure data in your environment as well as prepare or manage e-discovery requests that you receive. As Murphy said, “Companies are getting smarter about e-discovery. Because it's an inherently reactive process, they're going to expect that the people running their messaging systems are increasingly adept at getting this information and getting it quickly so that legal can start reviewing that





## B. K. Winstead

is a senior associate editor for *Windows IT Pro*, *SQL Server Pro*, and *SharePoint Pro*, specializing in messaging, mobility, and unified communications.

Email



Twitter



Blog



information within hours of a request as opposed to within weeks.” Having a good solution in place to find appropriate data before you get such a request could be crucial.

However, you might also find products that secure your environment, help enforce internal policies, prevent data leaks, and thereby reduce the chances that you’ll be asked to respond to e-discovery requests. As Foisey said, “Companies want to know who’s accessing the content, and specifically they want to be able to lock it down. When they find stuff, they want to be able to go in and say this is sensitive information in this location, and lock it down—regardless of what the Windows permissions say.” STEALTHbits offers [DLP Lite](#) to perform this level of protection.

Above all, don’t wait until you have some kind of snag to look for solutions to this growing data problem. You know data is accumulating, and the more data you have, the more likely you’re going to encounter some kind of trouble as a result of it. “In companies, IT often is sort of looked at as the support organization,” Murphy said. “I think there’s an opportunity for a lot of these IT pros, especially in the messaging world, to sort of be heroes in their organization by being proactive about how to handle issues that come up around e-discovery. There’s a real opportunity for them to have an impact on their organizations and even their careers.”

You have the power to take the fear out of e-discovery. Educate yourself and prepare. You might never receive an e-discovery request, but then again, do you want to take that chance? ■

—B. K. Winstead  
InstantDoc ID 142585



**Jason  
Bovberg**

Email



Twitter



# Stay Inside for the Summer!

## Product of the Month

Mowing your lawn can be such an onerous task, particularly if you've got systems to monitor at work! You probably wish that lawn would take care of itself as you attend to your IT duties. The new LB1500 SpyderEVO by


Kyodo America Industries cuts grass all by itself, so you can stay inside and type at your keyboard! The LB1500 SpyderEVO is a gas-free, robotic lawnmower that uses a four-prong blade to cut grass automatically without the risk of damaging flowerbeds, play areas, and so on.

To determine your yard's boundaries, the LB1500 SpyderEVO senses an included perimeter wire that you place around your lawn. Check out all of Kyodo's gadgets and mowers at the [Kyodo company website](#).






**Figure 1:** Long-winded



**Figure 2:** Well, now I'm nervous



**Figure 3:** Um, both, I guess

## USER MOMENT OF THE MONTH

I work in a legal office, and one day a senior lawyer called me to complain about what he called his “broken mouse.” When I arrived at his desk, he demonstrated the problem he was having: As he moved the mouse across his pad, the onscreen cursor moved jerkily. I took a few moments to open the mouse, remove the trackball, and blow out all the accumulated dirt and dust. I also wiped off the mouse pad. The mouse worked perfectly, and all was right with the world. But I got a call the next morning from the same lawyer. The mouse wasn’t working at all. Upon inspection, I found that the lawyer—sensing a design flaw in the mouse—had applied tape to the device’s underside, so that no more dirt would accumulate around the ball.

—Kevin Toth

Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.



**Submit**

Search our network of sites dedicated to hands-on technical information for IT professionals.

[www.windowssitpro.com](http://www.windowssitpro.com)

## Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.

[www.windowssitpro.com/go/forums](http://www.windowssitpro.com/go/forums)

## News

Check out the current news and information about Microsoft Windows technologies.

[www.windowssitpro.com/go/news](http://www.windowssitpro.com/go/news)

## EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- [Dev Pro UPDATE](#)
- [Exchange & Outlook UPDATE](#)
- [Security UPDATE](#)
- [SharePoint Pro UPDATE](#)
- [SQL Server Pro UPDATE](#)
- [Windows IT Pro UPDATE](#)
- [WinInfo Daily UPDATE](#)

## RELATED PRODUCTS

### Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.

[www.windowssitpro.com/go/vipsub](http://www.windowssitpro.com/go/vipsub)

### SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.

[www.sqlmag.com](http://www.sqlmag.com)

### Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.

[www.devproconnections.com](http://www.devproconnections.com)

### SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.

[www.sharepointpromag.com](http://www.sharepointpromag.com)

## Advertiser Directory

<b><u>Binary Tree</u></b> .....	159
<b><u>Centrify</u></b> .....	2
<b><u>Commvault Systems</u></b> .....	76
<b><u>eLearning Series</u></b> .....	115
<b><u>NetApp</u></b> .....	26
<b><u>Specops Software</u></b> .....	6
<b><u>Vision Solutions Inc.</u></b> .....	1
<b><u>Western Governors University</u></b> .....	68
<b><u>WinConnections Fall 2012 Event</u></b> .....	38
<b><u>Windows IT Pro Events Calendar</u></b> .....	75

## Vendor Directory

<b><u>Advanced Systems Concepts</u></b> .....	130
<b><u>Amazon Web Services</u></b> .....	118
<b><u>AppDev</u></b> .....	155
<b><u>AppRiver</u></b> .....	130
<b><u>AvePoint</u></b> .....	143
<b><u>Azaleos</u></b> .....	118
<b><u>BeyondTrust</u></b> .....	139
<b><u>Brocade</u></b> .....	127
<b><u>Browsium</u></b> .....	128
<b><u>CBT Nuggets</u></b> .....	156

<b><u>Citrix Systems</u></b> .....	132
<b><u>CloudShare</u></b> .....	118
<b><u>Connectria</u></b> .....	118
<b><u>Ericom</u></b> .....	129
<b><u>F5 Networks</u></b> .....	129
<b><u>FPWeb.net</u></b> .....	118
<b><u>HP</u></b> .....	8
<b><u>Idera</u></b> .....	143
<b><u>Jalasoft</u></b> .....	44
<b><u>K Alliance</u></b> .....	156
<b><u>Kyodo America</u></b> .....	165
<b><u>LearnKey</u></b> .....	156
<b><u>ManageEngine</u></b> .....	143
<b><u>Nokia</u></b> .....	131
<b><u>Quest Software</u></b> .....	143
<b><u>Rackspace</u></b> .....	118
<b><u>RIM</u></b> .....	73
<b><u>SkillSoft</u></b> .....	158
<b><u>STEALTHbits</u></b> .....	162
<b><u>TestOut</u></b> .....	158
<b><u>Thycotic Software</u></b> .....	128
<b><u>TrainSignal</u></b> .....	158
<b><u>VMware</u></b> .....	41
<b><u>XIntercept</u></b> .....	39